USE OF LANDMARKS TO IMPROVE SPATIAL LEARNING AND REVISITATION IN COMPUTER INTERFACES

A Thesis Submitted to the College of

Graduate Studies and Postdoctoral Studies

In Partial Fulfillment of the Requirements

for the Degree of Doctor of Philosophy

in the Department of Computer Science

University of Saskatchewan

Saskatoon, Canada

by

Md. Sami Uddin

Permission to Use

In presenting this thesis in partial fulfilment of the requirements for a Postgraduate degree from the University of Saskatchewan, I agree that the Libraries of this University may make it freely available for inspection. I further agree that permission for copying of this thesis in any manner, in whole or in part, for scholarly purposes may be granted by the professor or professors who supervised my thesis work or, in their absence, by the Head of the Department or the Dean of the College in which my thesis work was done. It is understood that any copying or publication or use of this thesis or parts thereof for financial gain shall not be allowed without my written permission. It is also understood that due recognition shall be given to me and to the University of Saskatchewan in any scholarly use which may be made of any material in my thesis.

Requests for permission to copy or to make other use of material in this thesis in whole or part should be addressed to:

Head of the Department of Computer Science
176 Thorvaldson Building
110 Science Place
University of Saskatchewan
Saskatoon, Saskatchewan S7N 5C9
Canada

Or

Dean
College of Graduate and Postdoctoral Studies
University of Saskatchewan
116 Thorvaldson Building, 110 Science Place
Saskatoon, Saskatchewan S7N 5C9
Canada

ABSTRACT

Efficient spatial location learning and remembering are just as important for two-dimensional Graphical User Interfaces (GUI) as they are for real environments where locations are revisited multiple times. Rapid spatial memory development in GUIs, however, can be difficult because these interfaces often lack adequate landmarks that have been predominantly used by people to learn and recall real-life locations. In the absence of sufficient landmarks in GUIs, artificially created visual objects (i.e., *artificial landmarks*) could be used as landmarks to support spatial memory development of spatial locations. In order to understand how spatial memory development occurs in GUIs and explore ways to assist users' efficient location learning and recalling in GUIs, I carried out five studies exploring the use of landmarks in GUIs – one study that investigated interfaces of four standard desktop applications: Microsoft Word, Facebook, Adobe Photoshop, and Adobe Reader, and other four that tested artificial landmarks augmented two prototype desktop GUIs against non-landmarked versions: command selection interfaces and linear document viewers; in addition, I tested landmarks' use in variants of these interfaces that varied in the number of command sets (small, medium, and large) and types of linear documents (textual and video). Results indicate that GUIs' existing features and design elements can be reliable landmarks in GUIs that provide spatial benefits similar to real environments. I also show that artificial landmarks can significantly improve spatial memory development of GUIs, allowing support for rapid spatial location learning and remembering in GUIs. Overall, this dissertation reveals that landmarks can be a valuable addition to graphical systems to improve the memorability and usability of GUIs.

ACKNOWLEDGMENTS

I want to express my gratitude to my mentor and supervisor, Dr. Carl Gutwin, for your guidance, inspiration, and input throughout my graduate career; nothing would have been possible without you. I am fortunate that I could work with you and learn directly from you for so many years. I thank you for making me who I am today.

I am also grateful to my advisory committee, Dr. Regan Mandryk, Dr. Cody Phillips, and Dr. Ehab Diab, for their expert opinions and critiques on my work that propelled me towards excellence.

A big thanks to the faculty and staff at the Department of Computer Science, particularly to the students and staff in the HCI Lab, for their help and friendship that created a lively and cooperative environment for carrying out research work, which made this long journey enjoyable and memorable.

I also thank all my friends in Saskatoon for welcoming me to the community and making this foreign place a home away from home for me. My special thanks to Md. Aminul Islam, Nazifa Azam Khan, Sakib Ahmed, Md. Aradot Ali, and Sowgat Ibne Mahmud, for those shared laughs, memories, and countless discussions at Tims – you brought warmth to my life during the long and cold Saskatoon winters!

Finally, I thank my mother, Mrs. Shamim Sultana, my father, Dr. Md. Nasir Uddin and my brother, Sayem Uddin, for providing me thorough support and courage to face challenges while living abroad, and last by not least, my wife, Sadia Tasnim Tuba, whom I met during my Ph.D. who never stops to amaze me through her love and support.

This dissertation is dedicated to my parents.

CONTENTS

LIST OF TABLES

LIST OF FIGURES

List of Abbreviations

1D	One Dimensional
2D	Two Dimensional
3D	Three Dimensional
AR	Augmented Reality
GUI	Graphical User Interface
HCI	Human-Computer Interaction
LTM	Long Term Memory
OS	Operating System
PDF	Portable Document Format
RM-ANOVA	Repeated Measures Analysis of Variance
RT	Reaction Time
SFT	Space-Filling Thumbnails
STM	Short Term Memory
TLX	Task Load Index
UI	User Interface
VR	Virtual Reality
WIMP	Windows, Icons, Menus and Pointers

CHAPTER 1

INTRODUCTION

Graphical User Interfaces (GUIs) are ubiquitous. We see them in everything from traditional desktop computers and laptops to popular handheld multi-touch smart devices and wearables. GUIs present graphical items as icons, menus, toolbars, and controllers (i.e., sliders or scrollbars) that usually appear at specific locations within a window; users interact with computers by visiting those locations. A critical part of being an expert with a GUI is knowing the locations of graphical items (e.g., tools/commands). As a result, experts can find graphical objects quickly and accurately [50]. However, novice GUI users typically rely on slow visual search [187] to find desired items since they do not know where the items are located. Spatial memory [23,170,208] is a powerful way for users to become experts with a GUI as it helps them remember item locations. If users of a GUI can easily remember item locations, they do not need to carry out a slow visual search. Spatial learning in the real world benefits significantly from landmarks in the environment. However, user interfaces often provide very few visual landmarks. This dissertation explores the use of *artificial landmarks* as a way to improve people's spatial memory and to support expertise development in computer interfaces.

1.1 PROBLEM

The general problem addressed in this dissertation is that people face difficulty in finding and re-finding spatial locations in GUIs that they have visited earlier. Part of this problem arises because GUIs often do not provide enough support for developing spatial memory of those locations. GUIs that employ the easy-to-use Windows, Icons, Menus, and Pointers (WIMP) paradigm present all tools and control mechanisms at particular locations in the interfaces that users must find and visit to complete tasks. Many of the tasks people carry out with computers are repetitive [86,87,206]; therefore, users need to repeatedly visit the same locations in GUIs. For example, command selection from menus or toolbars and navigation in linear media (e.g., a video or a PDF document) are two simple yet essential tasks that often involve repeatedly locating and revisiting spatial locations in user interfaces. One of the many goals of visual interface design is to support the user's transition from novice to expert [51,129]. However, revisiting spatial locations can be slow and erroneous, particularly for users who are just beginning to make that transition – that is, users who are just learning where the commands are – because GUIs often make it difficult to remember locations.

Spatial memory [23,170,208] is a powerful way to enable expert performance with user interfaces. One benefit of human memory with spatial locations is that people can retrieve the locations of already-visited items from memory and perform a quick revisitation. Spatial consistency in a GUI – the idea that interface elements do not change location – helps users develop and make use of spatial memory. Many research interfaces have exploited the efficiencies offered by stable spatial locations. For example, ListMaps [92] converts the items of a long list into a stable two-dimensional grid and allows users to develop spatial memory by showing all the items at once. Similarly, CommandMaps [183] allows access to large command vocabularies by flattening the traditional hierarchy and assigning each command to a unique spatial location in a display. Although spatially stable interfaces can enable high input efficiency when the user is an expert, the attainment of such expertise requires learning and efficient recall of item locations.

Spatial learning in the real world benefits significantly from landmarks because landmarks provide a stable reference frame for the locations of nearby objects. However, most GUIs do not provide adequate landmarks that can be used to support spatial memory development. Spatial learning and revisitation of locations occur at two different levels in GUIs: entire interfaces (e.g., command locations in menus and toolbars), and individual widgets (e.g., locations in a linear controller such as a scrollbar). In the GUI environment, implicit landmarks such as the corners of a computer screen can help users learn the items located near them; however, these landmarks are not part of the designer's intention and cover only a small portion of the interface. The majority of the areas in GUIs do not have such reliable landmarks to aid in spatial memory development. Therefore, learning and revisiting the spatial locations of commands can be slow and difficult because users must rely on absolute spatial memory, which can only get users to the general vicinity of a command. The lack of landmarks makes locating a command in such GUIs time-consuming. Similarly, hindrances in spatial learning can be seen in linear documents (e.g., videos or PDFs) because 1D control widgets (e.g., scrollbars or sliders) have no clear visual landmarks, and all locations on the widgets look similar.

This dissertation, therefore, addresses the problem that ***GUIs do not provide adequate landmarks to support spatial learning*** and aims to develop effective ways to assist users in learning and recalling spatial locations in graphical user interfaces quickly.

1.2 MOTIVATION

One goal of designing GUIs is to support users in becoming experts. It is important to improve people's level of expertise with GUIs because if a GUI can make users experts quickly, it will improve efficiency and user satisfaction. Therefore, it is essential to investigate how GUIs can support users' spatial learning. GUIs are the most common user interface mode available in modern computing devices. Before GUIs, however, early computers used command line-based interfaces that required users to type textual commands. Although experts could eventually type commands quickly and accurately in a command-line-based interface, the development of such expertise was slow and only possible after extensive learning. In contrast, GUIs present graphical elements, such as windows, menus, commands, links, and visual controllers, that users can easily manipulate with minimal training. Although GUIs have liberated users from learning and remembering textual commands, efficient interaction with a command in a GUI depends on how quickly users can find its location.

Since GUIs present commands visually, a user can find the desired command after visually searching for it, but finding one command out of many could be difficult for a novice user of that GUI. However, as the Power-Law-of-Practice [77] suggests, after enough practice, expert users of a GUI can find commands quickly and accurately because they have developed a memory of those commands' locations on the screen. A key differentiating factor between novices and experts is *knowledge of a command's location*. If we can understand how users develop spatial memory of commands, we can exploit that information to design improved GUIs to support users' rapid expertise development in GUIs. Therefore, it is essential to investigate what contributes to users developing expertise with a GUI in terms of learning locations of its graphical elements (e.g., commands).

The capability of learning and recalling an object's location in humans is primarily governed by spatial memory [23,170,208]. Since most human interactions with GUIs, such as selecting commands while writing texts or editing graphics and (re)visiting episodes while browsing PDFs or watching videos, involve finding locations in GUIs, spatial memory can enable users to perform these actions efficiently as experts with these GUIs can simply recall locations from memory, without visually looking for the location in the interface. Spatially consistent menu designs [95,183] that always show graphical items (e.g., commands) in spatially stable locations on a screen can enable spatial memory development. Early works with spatially stable menus such as ListMaps [92], CommandMaps [183], and linear controller-based Footprints Scrollbar [7] showed spatial memory's performance advantages over slow visual search-based methods. However, a large number of items in an interface (e.g., hundreds of commands in the MS Word application) and similar-looking locations in 1D control widgets often hinder the development of the spatial memory of locations and, in turn, make remembering locations accurately challenging.

Landmarks present in an environment contribute to the development of spatial memory of locations in real life by providing strong reference points for the objects in the environment [15,143]. People can quickly remember the locations of objects in a known area by relying on landmarks in the area. However, such landmarks are uncommon in computer interfaces. The graphical interfaces of computers often provide no landmarks, causing the development of spatial memory to slow. If GUIs had better and stronger landmarks in interfaces, they could expedite location learning and, in turn, could facilitate quick and accurate recall of those locations. Therefore, it is crucial to understand whether landmarks support spatial memory development in GUIs, what can act as landmarks in GUIs and determine effective ways to integrate landmarks in interfaces to aid users in quickly developing spatial memory of locations.

1.3 SOLUTION

As a solution to the problem of GUIs having inadequate landmarks to support spatial learning, this dissertation investigates the use of *artificial landmarks* in computer interfaces. Artificial landmarks are artificially created digital artifacts or elements, such as solid colour blocks, images, and unique or random icons that can be used as landmarks in graphical user interfaces. In the absence of strong and natural landmarks in GUIs, these artificial landmarks can help users develop spatial memory of digital contents quickly and remember those locations efficiently. Although artificial landmarks are functionally indifferent to natural landmarks, in this dissertation, the term *'artificial landmark'* has been consciously used to indicate the proposed novel landmarks (see Chapters 4 and 5) used in digital spaces. The existing features in the GUI environment, such as the corners of a screen, are simply referred to as landmarks. However, for simplicity, we can treat them together as just 'landmarks.'

Landmarks are easily distinguishable and permanent objects in an environment that can help people remember or identify nearby objects [15,143]. We use landmarks to remember and navigate through different real-life locations. For example, a person may say, "meet me by the bench beside the pine tree." Here, the pine tree is a landmark that references a specific bench near it. A few landmarks, such as the corners of a device, are currently available in digital interfaces. Several techniques have already utilized these landmarks to organize interface items to develop users' spatial memory of commands and improve location recall performance in small handheld devices where the number of commands is limited [95,132,191]. However, these natural landmarks become inadequate for commands in relatively large devices, such as desktops, particularly in the middle region of desktop interfaces consisting of many commands. As a result, learning and remembering spatial locations in GUIs become difficult.

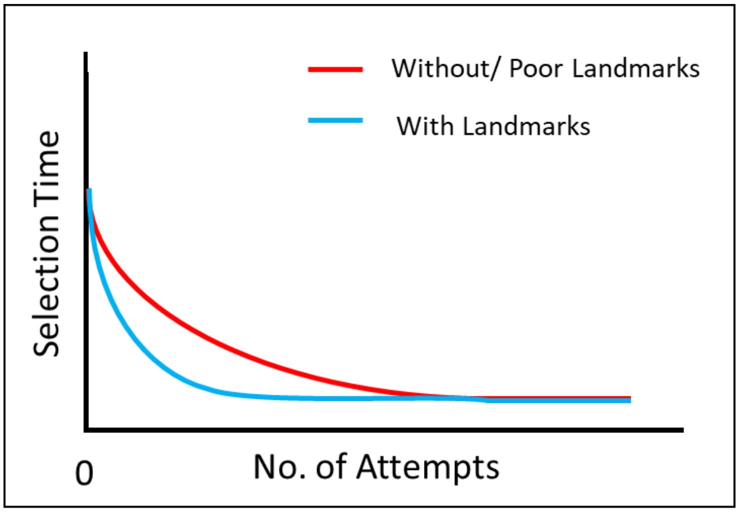

Figure 1.1: Hypothetical Power-Law-of-Practice curves for interfaces with landmarks and without or having poor landmarks.

The learning and remembering of spatial locations in GUIs can be described by the Power-Law-of-Practice [77], which indicates that with sufficient practice, users of a GUI can develop a memory of spatial locations. The absence of adequate landmarks in these GUIs, however, makes the process of spatial location learning slow (see Figure 1.1: the red line). As a solution, artificially planted objects in GUIs can be used as landmarks. Similar to the landmarks available in real life, these artificial landmarks can also provide strong reference points for efficient learning and remembering spatial locations in GUIs. As a result, the time required to learn spatial locations can be reduced significantly (illustrated in Figure 1.1: the blue line), and novices can quickly become experts with a computer interface.

Therefore, the primary goal of this dissertation is to investigate the effects of using *landmarks* in graphical interfaces to develop users' spatial memory of interfaces. However, before augmenting 2D computer interfaces with artificial landmarks to see how they affect the performance of learning and revisiting spatial locations in GUIs, we need first to understand how spatial memory development takes place in standard GUIs. So, in this dissertation, I focus on three secondary goals concerning spatial memory development in GUIs. First, I investigate how people learn spatial locations in standard computer interfaces that people regularly use and whether landmarks contribute to that learning process. Second, I explore the use of artificial landmarks at the level of an entire GUI, and as a representative, I consider a command selection interface. Last, I examine how artificial landmarks perform at the level of a widget in a GUI, such as linear document viewers' controllers (i.e., sliders or scrollbars).

1.4 SUMMARY OF STUDIES

In order to accomplish my dissertation research goals, I carried out five studies[1] that investigated commercially available standard GUIs and evaluated prototype interfaces augmented with different artificial landmarks. In the first study, I investigated interfaces of four popular commercial desktop applications to understand how people develop spatial memory of those GUIs and whether any existing landmarks in GUIs contribute to spatial memory development. In the second study, I used grey colour blocks and an image as the menu backdrop for landmarks in GUIs to investigate if those artificial landmarks could improve learning the locations of commands in a command selection interface having a small number of commands. In the third and fourth studies, I investigated the effects of landmarks in medium and large command-set sizes. The landmarks and the command selection interfaces used in these two studies were identical to the second study. Finally, in the fifth study, however, I explored two new artificial landmarks (i.e., random icons and thumbnails) and tested how they supported developing spatial memory in linear documents.

I arranged the five research studies into three manuscripts. Each of these three manuscripts focuses on specific aspects of my overall goal of understanding users' spatial memory development in GUIs and examining the role landmarks play in improving it. Manuscript A[2] presents the first study focusing on spatial memory development in commercially available desktop applications. It looks if landmarks are present in those GUIs and provide any benefit in spatial learning and recall.

---

[1] All studies and manuscripts presented in this dissertation are products of a collective effort and completed with support from my colleagues. In order to highlight my efforts and contributions in carrying out the research in cooperation with my supervisor, contextualize my research as part of my dissertation, and to differentiate collective contributions, I will refer to "we" in the manuscript sections and "I" in rest of the dissertation sections.

[2] **Manuscript A:** Md. Sami Uddin, and Carl Gutwin. 2021. The Image of the Interface: How People Use Landmarks to Develop Spatial Memory of Commands in Graphical Interfaces. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, ACM, New York, NY, USA, 17 Pages.

Manuscript B[3] presents the three subsequent studies (i.e., second, third, and fourth) investigating the effects of artificial landmarks in learning and remembering the locations of commands in GUIs, where the total number of commands is progressively increased. Finally, Manuscript C[4] investigates the effects of artificial landmarks in developing spatial memory of linear documents with two different linear document controllers: a slider and a scrollbar (i.e., the fifth study). Manuscripts[5] A, B, and C have been published in leading HCI venues, and Manuscript A received an ***Honorable Mention Award*** (top 5% paper) at CHI 2021!

1.5 CONTRIBUTIONS

This Ph.D. dissertation makes several contributions to HCI.

First, it reveals that people develop spatial images of the frequently used interfaces in their minds and provides evidence that people use these cognitive images to recall the locations of previously visited commands in the interfaces.

Second, my dissertation presents and demonstrates a methodology to elicit the images of interfaces from users' minds. Analyses of the cognitive images of four commercial GUIs revealed new information that standard GUIs have four types of landmarks upon which people strongly rely to learn and recall the locations of commands.

Third, it empirically demonstrates for the first time how the use of *artificial landmarks* in GUIs can aid in better location learning of graphical elements (e.g., commands/tools and episodes in

---

[3] **Manuscript B:** Md. Sami Uddin, Carl Gutwin, and Andy Cockburn. 2017. The Effects of Artificial Landmarks on Learning and Performance in Spatial-Memory Interfaces. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, ACM, New York, NY, USA, 3843–3855.

[4] **Manuscript C:** Md. Sami Uddin, Carl Gutwin, and Alix Goguey. 2017. Using Artificial Landmarks to Improve Revisitation Performance and Spatial Learning in Linear Control Widgets. In *Proceedings of the 5th Symposium on Spatial User Interaction*, ACM, New York, NY, USA, 48–57.

[5] The order of the three manuscripts presented in this dissertation differs from their publication chronology. Although Manuscript A was the latest among the three, it was presented first in the dissertation to construct a better research argument.

linear documents such as videos and PDFs) and allow efficient revisitation of those locations in the future in two different contexts: command selections and linear document revisitations.

Fourth, my dissertation demonstrates four types of artificial landmarks (colour blocks, an image as a menu-backdrop, random icons, and thumbnails from documents) that can be used to augment computer interfaces. It also demonstrates how they perform in different interfaces varying in sizes and types.

Last, it provides guidelines for designers of future interfaces to design more memorable and efficient GUIs with the help of artificial landmarks.

1.6 DISSERTATION OUTLINE

This dissertation is organized into several chapters. Chapter Two summarizes the foundations for this research. It provides a brief overview of human memory and topics related to learning and recalling locations. It also reviews related research and techniques for understanding and supporting spatial memory development in graphical interfaces.

Chapters Three to Five present Manuscripts A, B, and C – each chapter begins with an introduction to the studies included in the manuscript, then discusses the research problem addressed in the manuscript, the solution to that problem and the approach I followed to reach the solution. The chapter then presents the actual manuscript, followed by a brief summary of findings and contributions. Each chapter ends with a section highlighting the relevance of the presented work in relation to the overall dissertation.

Chapter Six discusses the findings from the three manuscripts and how they fit in the context of my overall dissertation research and literature. Then I discuss some high-level insights that I came across over the course of this dissertation and shed light on potential future work. Finally, it concludes by summarizing the main contributions of this dissertation. The supplementary materials used to carry out studies presented in this dissertation, such as consent forms, questionnaires and washed-out interfaces, are included in the appendices.

CHAPTER 2

BACKGROUND ON LOCATION LEARNING AND RECALL

This chapter provides a general overview of the theories and models available in the literature related to human location learning and recall. Since this dissertation focuses on learning and remembering locations of graphical objects in computer applications, we need to understand the concepts underlying location learning in GUIs, including cognitive factors contributing to location learning and recall as well as the spatial design of graphical interfaces. Therefore, this chapter summarizes the current understanding of cognitive factors associated with learning locations in real life, including the interface design practices supporting location learning in GUIs.

2.1 MEMORY OF LOCATION

In this section, I present an overview of cognitive models and literature on memory, particularly memory related to learning and remembering the locations of objects. Spatial memory enables people to learn and recall the locations of objects present in a space. However, spatial memory is not a single memory unit; instead, several components of the human memory system are involved in remembering locations. Therefore, after introducing spatial memory, I present a brief review of the memory systems involved in learning and remembering objects' locations.

2.1.1 Introduction to Spatial Memory

Spatial memory is one part of human cognition that allows people to record spatial information for objects in an environment. People use this memory to store information about the object's location, including spatial relations among objects in both short and long terms. Spatial memory enables people to remember where an object is in an area. Human spatial memory also involves "the manipulation and orientation of single objects (e.g., mental rotation and knowing the location of an object relative to a reference point such as the body) and spatial orientation, which is orienting in large-scale space (e.g., spatial navigation and wayfinding)" ([111], p. 86). Taken together, these capabilities enable people to "manipulate, recall, and navigate through space" ([111], p. 86), be it in the physical world or digital spaces. In the following sections, I highlight the spatial tasks people typically carry out on computers that require learning object locations, and then I discuss object-location memory.

2.1.1.1 Spatial Tasks in Computers

Remembering the locations of graphical elements (e.g., commands) and navigating through them are two categories of common spatial tasks in GUIs, particularly in static computer displays. This dissertation distinguishes between these two spatial tasks while discussing the related literature. Although both tasks are spatial in nature, the two task types can be different depending on an observer's operational perspective in a GUI. For instance, Maguire et al. [145] indicated that recalling spatial locations on a display is usually executed from an aerial viewpoint, while navigation occurs from a viewer-centred perspective. That is, a user can see the entire space when recalling locations in computer displays; in contrast, only a portion of the space is visible to the user while navigating. As a result, researchers such as Hegarty et al. [98] have found only a weak correlation between the two spatial tasks. Neuroscience researchers have also acknowledged the difference between these two spatial tasks: "navigation is not the same as table-top tests of spatial memory [...] direct inferences cannot be made about one from the other" ([145], p. 171). There is, however, one area – interaction in large-scale environments – where remembering spatial locations and navigation can occur together. Siegel and White's model of developing spatial knowledge in a large space indicated an overlap of these two spatial tasks [195].

Although navigation can be crucial to spatial tasks in computers, such as wayfinding in virtual environments [64,65], most tasks in traditional desktop interfaces require users to remember the locations of graphical objects (e.g., commands) that appear at particular locations on screens.

Therefore, I focus primarily on the memory responsible for learning and remembering the locations of objects - object-location memory.

2.1.1.2 Object-Location Memory

This section focuses on the practical steps involved in learning and remembering the locations of objects in general. Object-location memory [18,172] is a critical part of spatial memory that allows people to remember an object's position in an environment. Researchers have investigated various aspects of object-location memory. For example, Hasher and Zacks [97] argued that the operation of the object-location memory takes place automatically. However, there is some evidence [171] arguing that the process of object-location memory requires explicit attention. Postma et al. [171] proposed a model of object-location memory consisting of three practical steps: object processing, spatial-location processing, and binding of the objects to locations.

a) *Object processing.* According to the object-location memory model, remembering the location of an object does not begin with processing spatial information; instead, it starts with the "representation or description of where things are in space, independent from how and in which order the observer wants to attend to these locations" ([171], p. 143). In other words, recalling an object's location begins with a standard object recognition process involving recognizing or identifying an object by inspecting its visual properties.

b) *Spatial-location processing.* After processing an object's visual appearance (i.e., recognizing the object), its location information is processed. Spatial locations can be represented in categorical representation and coordinate representation. Categorical representations "refer to relative spatial relations, such as remembering that your cup is on the right of the computer" ([18], p. 250) and describe objects using certain relations (e.g., left/right, above/beneath, near/far), without specifying exact location information. Coordinate representations specify locations using units in a reference frame - they "contain fine-grained, metric information, which can be used to guide actions, particularly when visual information is not at hand or insufficient" ([18], p. 250). Research also indicates that people are most likely to rely on categorical representation when several objects are present in an object-location memory task [6].

c) *Binding objects to locations.* The third step of this object-location memory model is responsible for integrating the information gathered in the previous two steps. The process of binding objects to locations – which can involve both effortful and automatic processes [171] – “requires connecting object identity information to either an exact position [coordinate representation] or a relative position [categorical representation]” ([172], p. 1340).

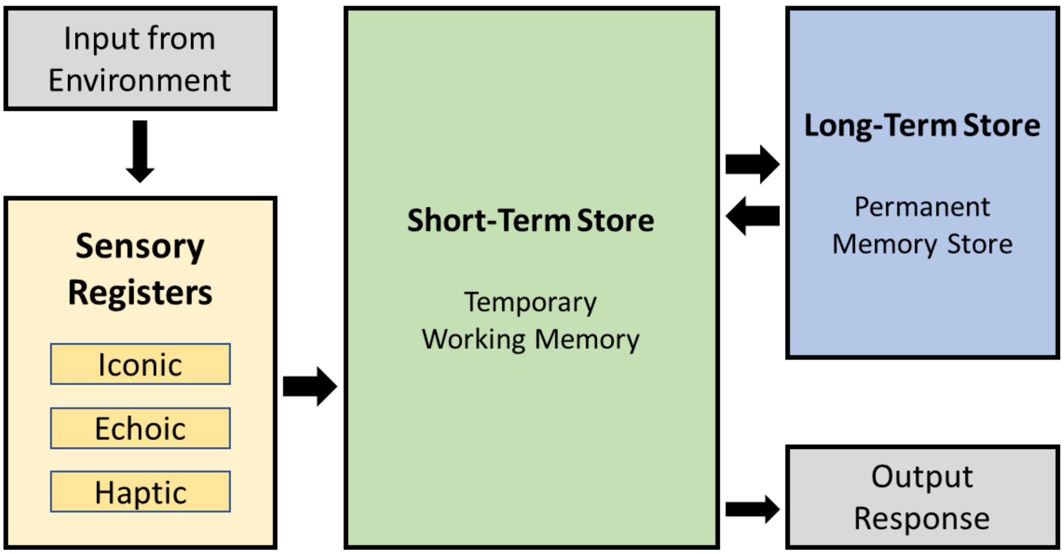

Figure 2.1: Atkinson and Shiffrin’s memory model (adapted from Atkinson et al. [19], p. 93).

2.1.2 A Brief Overview of Human Memory

Human memory [26,226] is a powerful cognitive process responsible for acquiring information, encoding it, storing it, and later retrieving it when needed. It has numerous implications in our lives, and it governs how we experience the world we are living in, from recalling meaningful incidents to empowering us to carry out tasks. Human memory does not operate in isolation; the brain is not only responsible for storing information but also for processing and acting on that information. To understand how people learn and remember the locations of objects in GUIs, we need to understand the components that make up the human memory system. Atkinson and Shiffrin [19] proposed a memory model (see Figure 2.1) consisting of three main memory types: sensory memory, short-term memory and long-term memory. The following sections review their memory model.

2.1.2.1 Sensory Memory

Sensory memory acquires information from an environment through sensory organs (e.g., eyesight or touch) and stores it for a very short period. It acts as a memory buffer that can briefly hold a large amount of data [196], which is erased unless we pay attention to it, which forwards it to short-term memory for further processing. Three sensory memory types are discussed in the literature: iconic memory, echoic memory, and haptic memory.

Iconic Memory. Iconic memory refers to the *visual sensory memory* that stores the cognitive representation of visual stimuli, which is information collected through sight. Sperling first introduced iconic memory in 1960 [201], and later Averbach et al. empirically validated its existence [20]. Iconic memory is particularly relevant to this dissertation because users interact with the locations of graphical items in desktop computer displays through visual sensors (i.e., their eyes) and store that location information in their iconic memory. Therefore, the development of spatial memory for graphical items begins by storing visual details on those items in iconic memory.

Researchers have identified three features of iconic memory [56,201]. First, iconic memory has a large storage capacity. Sperling empirically demonstrated that subjects could recall about 80% of 12 alphanumeric characters shown to them [201]. Second, despite the large capacity, visual information stored as iconic memory degrades rapidly [20], which indicates that iconic memory has a short life span (approximately one second). However, there is some evidence that some people's iconic memory can last longer, which is often treated as eidetic imagery or photographic memory [14]. Last, Sperling considered iconic memory as "pre-categorical" in nature as it is stored as raw visual data without meaning.

Echoic Memory. Echoic memory is the auditory sensory memory that stores information collected through auditory sensors– that is, audio data that people can hear. In contrast to iconic memory, where we can examine visual stimuli multiple times, we cannot scan auditory stimuli many times as the stimuli enter the echoic memory only once. Therefore, echoic memory has the capacity to store data slightly longer than visual memory – about four seconds [60].

Haptic Memory. Haptic memory, also known as tactile memory, is a specific form of sensory memory that stores information collected through touch. Although our whole body is capable of receiving haptic information, the perception of haptic information "takes place largely by inspection of shapes by the palm and fingers of the hand as they move over the surface of an object" ([109], p. 589). Bliss et al. first investigated haptic memory and revealed that haptic memory can store roughly five items [34] and has a short life span analogous to visual sensory memory [201]. Although there are differences between the haptic and visual memory systems, studies suggest an object's information can be stored and shared between these memory systems [10,70].

2.1.2.2 Short-Term Memory

Short-term memory (STM), as the name suggests, enables people to temporarily store a small amount of information for a short duration, usually ranging from a few seconds to one minute [24]. Atkinson and Shiffrin's model established the existence of a short-term storage facility for data [19]. As depicted in Figure 2.1, STM acts as a processing unit that receives information from both sensory and long-term storage, and prepares reasonable responses. The sensory registers continuously collect information from the environment and temporarily store it in sensory memory, but the data is erased unless it is attended to and forwarded to STM for further processing. Later, the processed data is sent to long-term memory for permanent storage. Atkinson and Shiffrin's model considers STM as the central component of the memory system that is not only responsible for temporarily retaining and processing information but also for collecting new data and making it available for future use.

Further research on memory systems provided more comprehensive explanations of human memory [59,61]. A widely accepted memory model is Baddeley and Hitch's [27] structural model, which proposed a component-based *working memory* model.

Working Memory. Although *working memory* is described as a kind of memory, it is not memory storage like STM; instead, it is a combination of several memory systems. Baddeley identified working memory as "a brain system that provides temporary storage and manipulation of the information necessary for such complex cognitive tasks as language comprehension, learning and reasoning" ([21], p. 556).

According to Baddeley and Hitch's model, working memory consists of four subsystems [22,27]. The most prominent among the four is the *central executive*, which primarily controls attention [28] and coordinates the other three subsystems. First, the *phonological loop* assures the recognition of verbal information. Second, the *visuospatial sketchpad* stores visual and spatial information. Last, the *episodic buffer* provides a temporal representation of information stored in other subsystems.

The visuospatial sketchpad is particularly relevant to this research. It primarily generates and retains a visuospatial depiction of the visual world [142], usually acquired from visual sensory memory. As a result, the sketchpad enables people to perform visual and spatial tasks, such as driving a car while visualizing a route in memory or remembering visual images of objects or people's faces. It has also been shown that the visuospatial sketchpad helps people keep track of their location with reference to other objects while moving around an area [170].

The sketchpad is also responsible for retrieving and manipulating visual and spatial information stored in long-term memory. For example, one can recall the spatial layout of one's house from memory and visualize it (using the visuospatial sketchpad) while remembering the number of windows at the front of the house. The sketchpad can also help carry out cognitive tasks, particularly where alternative solutions are present. For example, research has shown that people can successfully recall long lists of directions by encoding information into the visuospatial sketchpad [40].

2.1.2.3 Long-Term Memory

As the name suggests, long-term memory (LTM) refers to storing information in memory for an extended time. Even though people may forget some of the information they have learned over time, some information can be saved in memory for a lifetime. LTM is the final stage of Atkinson and Shiffrin's model (Figure 2.1), where the information processed in STM is sent to be stored with the possibility of retrieval when needed. In addition to the long duration of memory, LTM differs from STM in that it offers a nearly infinite storage capacity for information, thus creating a platform for people to carry out many kinds of learning [138], including the location of graphical elements.

Like STM, LTM consists of several memory systems [26], which usually fall under two broad categories: declarative and non-declarative memories. Declarative memories are the information that people consciously remember (e.g., events and facts), while non-declarative memories are recalled subconsciously and effortlessly (e.g., skills such as swimming or riding a bike). However, spatial memory – the memory responsible for remembering the locations of objects – does not seem to be an explicit part of any LTM model; instead, multiple components from both declarative and non-declarative memory can contribute to spatial location learning and recall. For instance, locations can be recalled as facts – a declarative memory component (e.g., "my wallet is in my front-right pocket") – or as procedures – a non-declarative component (e.g., developing 'motor memory' by reaching for the wallet that is usually kept in the front-right pocket). I outline the components of LTM in the following sections.

Declarative Memory. Declarative memory, often called *explicit memory*, refers to information or memories that require conscious thought to learn and recall [221]. For example, when people try to remember information about specific events, the name of a person, or semantic facts about an environment, that information is retained in declarative memory. Declarative memory can be divided into two categories. First, *episodic memory* refers to storing and retrieving information related to one's personal life and day-to-day experiences, including information about events (e.g., place, date, and time). It is believed the recollection of these episodic memories can enable people to revisit the events mentally [214]. Second, s*emantic memory* stores factual information "that has been learned, but for which specific 'time and place' information about the source of the original experience is typically not known" ([68], p.87). It includes information about general facts, concepts, categories, historical events, and names, such as the team that Lionel Messi plays for or the city where the CN Tower is located.

Non-declarative Memory. Non-declarative memory, often called *implicit memory*, refers to all the memories that people can recall subconsciously. This type of memory also includes abilities and skills that allow people to carry out certain tasks by remembering the necessary processes without intentionally thinking about them [113]. Non-declarative memory consists of three types of memory: procedural, associative, and priming.

a) *Procedural memory* involves motor movements and executive skills required to carry out specific tasks [43]. This memory can perform at a subconscious level; therefore, the skills and procedures needed for completing tasks – even if they involve complex motor and instinctual activities – can be automatically recollected without conscious attention. It enables us to carry out day-to-day activities, such as riding a bike, swimming, or walking.

b) *Associative memory* refers to the specific memory that can be recollected by forming an unconscious association with other information. This memory development, therefore, involves associating responses to environmental factors and stimuli [13]. Later, when people encounter similar stimuli, they subconsciously recall the responses or experiences related to the stimuli.

c) *Priming* is another type of non-declarative memory that refers to changes in behaviour because of earlier experiences that may have occurred repetitively. It is the effect where "exposure to certain stimuli influences the response given to stimuli presented later" ([43], p.10). As a result, priming can even influence a person's thoughts flow by activating certain associations or facts in memory. For example, when asked to name an animal starting with the letter "C," many people might say "Cat" because of its popularity, but a person might say "Cow" because they have a previous association. Such recollection of memory occurs naturally and implicitly.

2.2 THE PROCESS OF LEARNING LOCATIONS

This section describes how location learning occurs in humans. Specifically, it discusses location learning from the learner's cognitive perspective and presents a model for processing information to produce responses to given stimuli. The rest of the section focuses on the process of learning the locations of real-world objects.

2.2.1 Spatial Knowledge Acquisition

Since people live in an environment and perform activities in it, "they perceive surrounding space and acquire knowledge about it" ([107], p. 94). This knowledge "includes the identities of places and landmarks, the patterns of path connections between places, distances, and directions between places" ([107], p. 94). People use that spatial knowledge to orient themselves within the environment and perform successful navigation. Researchers have looked at how spatial knowledge is acquired [155,170,208]. Siegel and White introduced a theoretical framework suggesting that people go through three consecutive knowledge stages (landmark, route, and survey) when acquiring spatial knowledge [195]. Although the framework has mixed support from empirical results, its influence on psychology (e.g., [84,209]) and geography (e.g., [139]) has led to it being called the *dominant framework* [155].

a) *Landmark knowledge*. After entering a new environment, people subconsciously start learning about available objects [97]. This automatic observation-based learning (often known as incidental learning) [149,227] helps to form landmark knowledge [74]. Ishikawa et al. [107] defined landmark knowledge as "knowledge about the identities of discrete objects or scenes that are salient and recognizable in the environment" (p. 94). Landmarks are easily identifiable permanent objects in an environment that can be distinguished from nearby objects [143]. However, Siegel et al. [195] treated all "unique configurations of perceptual events" (i.e., visual patterns) as landmarks that can represent particular geographic locations (e.g., all the nodes in Figure 2.2: left). Landmarks are the points in an environment from which a person starts to develop their understanding of the environment's spatial representation. Section 2.3.2 presents a detailed review of landmarks.

b) *Route knowledge.* Once people become familiar with the context of the environment, they begin to build route knowledge [209,224]. As shown in Figure 2.2: middle, route knowledge refers to "a sequence of paths where nodes correspond to the main landmarks previously memorized" ([174], p. 4). Siegel et al. [195] indicated that "if one knows at the beginning of a 'journey' that one is going to see a particular landmark (or an ordered sequence of landmarks), one has a route" (p. 24). Therefore, a route can be treated "as a sequence of objects and events" ([224], p. 44).

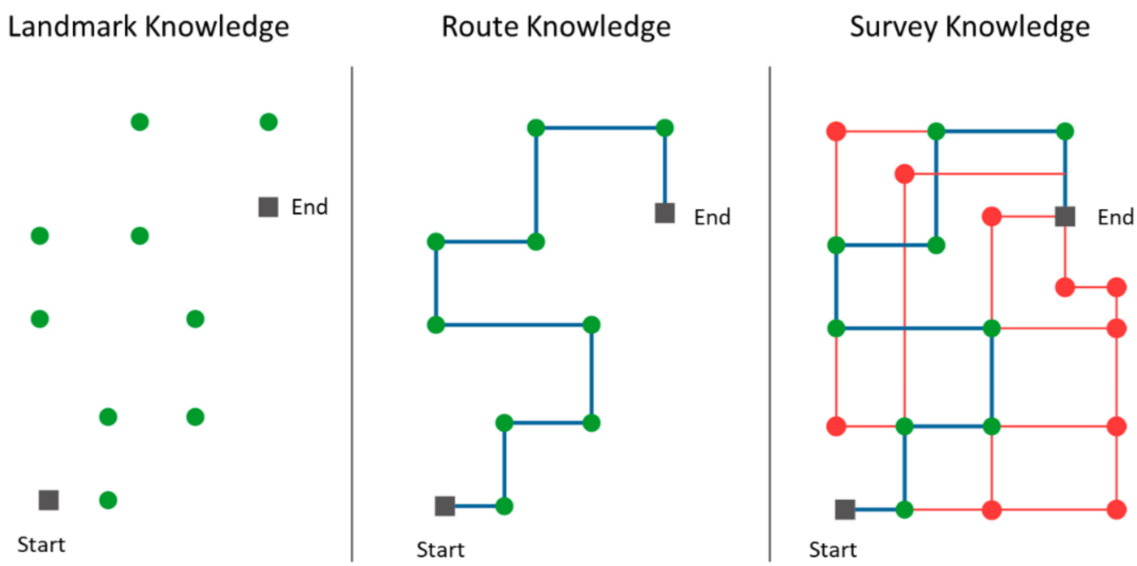

Figure 2.2: Representation of Landmark-Route-Survey knowledge (adapted from [174]).

c) *Survey knowledge.* The last and most complicated stage of spatial knowledge is survey knowledge. At this stage of spatial knowledge, a person has a map-like two-dimensional representation of the environment's entire configuration, including knowledge of all possible relationships (i.e., distances and directions) among the landmarks (see Figure 2.2: right). In order to acquire survey knowledge, "places and routes learned during separate travel experiences are integrated and interrelated with each other in a common frame of reference" ([107], p. 97). The acquired survey knowledge serves as a mental map of the items present in the environment – also called a *cognitive image* [212]. This cognitive image enables people to recall an item or perform navigation in the environment by providing necessary spatial information.

Researchers have raised questions about certain aspects of the dominant framework [107,155]. For example, the framework suggests that overall acquisition of spatial knowledge is slow, and "landmark knowledge is a necessary prerequisite for route knowledge, which in turn is a necessary prerequisite for survey knowledge." ([107], p. 95). That is, the development of spatial knowledge follows a sequential or hierarchical process, and the attainment of survey knowledge requires extensive experience with an environment [163]. However, research has found contradictory results indicating that even with limited exposure to a new area, people can take shortcuts and predict routes between locations – tasks that usually require survey knowledge of an area [125,140]. As a solution, Montello [155] proposed a *continuous framework*, which suggests a continuous development of spatial knowledge (i.e., the three stages of knowledge acquisition occur in parallel) [155], in contrast to the dominant framework's serial acquisition process. Although the continuous framework differs slightly from the dominant one, together, they enrich our understanding of how people develop spatial knowledge.

2.2.2 Power Law of Practice

Researchers have studied the time required to learn new information, such as the locations of items in an environment. These studies have shown that the *power law of practice* describes the relationship between mean response time and practice as a nonlinear function [75]. Figure 2.3 depicts the power law of practice. According to the law, the time required for a learner to complete a task decreases with practice (i.e., the number of attempts), and performance follows the shape of a power law (Figure 2.3). This type of graph is also known as a *learning curve* [4,44].

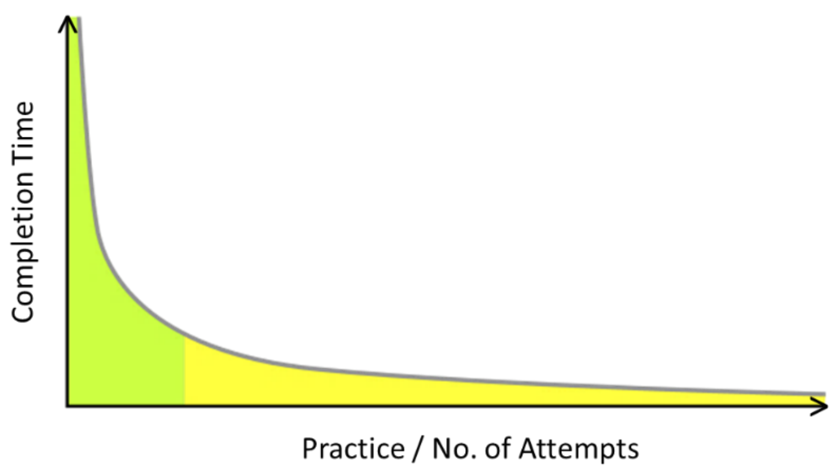

Figure 2.3: Graph of the Power-Law-of-Practice. The green area indicates the dramatic improvement in performance; yellow denotes the area with a slow rate of performance improvement.

As shown in Figure 2.3, performance dramatically improves during the early stages of learning (green area in the chart), but the improvement rate slows in later stages (yellow area) and eventually reaches a plateau. It is believed that learners achieve a *performance ceiling* (i.e., the maximum level of performance) [188] when they reach that learning plateau because of the learners' psychological restrictions (e.g., lack of motivation [119]) or a system's physical restrictions (e.g., a fixed cycle time of a machine [62]). There is also a *performance floor* that indicates a lower limit on performance (i.e., the maximum time that would be needed to complete a task for a first-time user) [188].

Researchers such as Fitts and Posner [77], Newell et al. [159], and Anderson [12] introduced the law to HCI. Card et al. [44] studied one user's learning performance over thousands of trials and indicated that user performance in computer interfaces could be described with the power law of practice. Selecting a command from a menu, which involves specifying its location, also follows the law. In the beginning, a user requires additional time to locate an item visually, but once the location is learned through practice, a user easily recalls its location from memory to perform a quick selection. Ahlstrom et al. [4], Cockburn et al. [48], and Scarr et al. [186] observed users' command selection behaviours in several types of menus and modelled users' performance with the power law of practice to describe how users transition from novices to experts in GUIs.

2.2.3 Developing Spatial Skill

Learning and remembering objects' locations can be developed through practice. Although the development of a skill is a continuous process, researchers have identified several stages that a learner has to go through [77,199]. One widely accepted model of skill development is the model proposed by Fitts and Posner [77,188,213], consisting of three phases: cognitive, associative, and autonomous (see Figure 2.4).

2.2.3.1 Cognitive Stage

The cognitive stage of skill development refers to the initial stage of performance, where people begin to understand and learn the process of carrying out a new task [77]. The process starts with the user becoming familiar with the required actions to achieve the desired results. Although the process of performing a task is often based on instructions (i.e., declarative knowledge) [12,123], people also follow an exploratory approach where they learn by attempting the task and making mistakes.

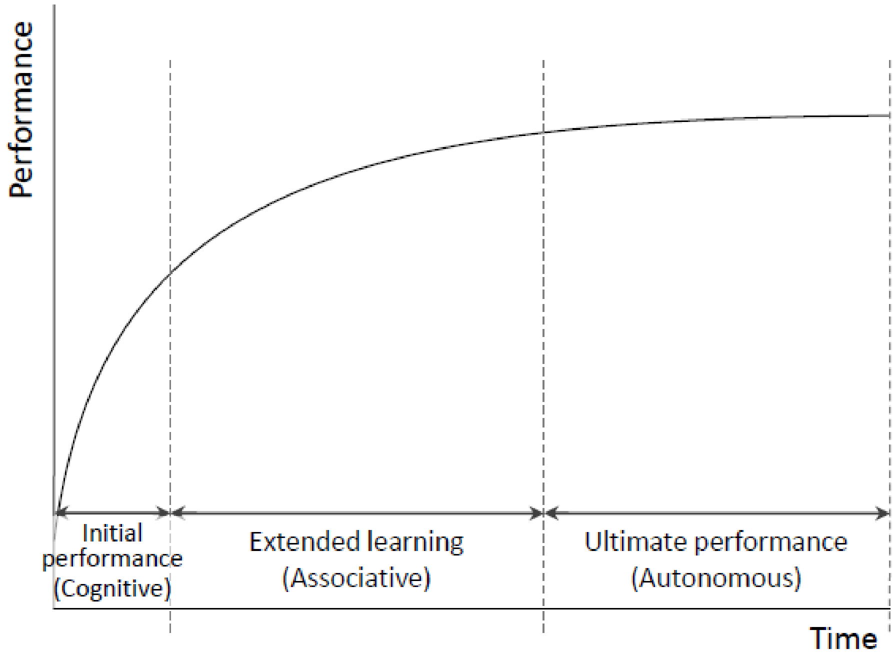

Figure 2.4: Three stages of skill development [187].

As shown in Figure 2.4, performance at the beginning of the cognitive stage is comparatively low because people have limited prior knowledge of the task. People usually deploy various trial-and-error strategies [188] and continuously monitor the outcomes to determine an effective way to carry out the task [204]. Since people are just beginning their learning in this stage, Fitts and Posner [77] suggested that it is the best stage to provide support (e.g., cues or instructions to respond in a certain way) to learners. The cognitive stage is where people are more likely to incorporate those supports in their learning process to come up with efficient strategies to complete a task.

Upon repeating the learned strategies, learners start to perform significantly better in the latter part of the cognitive stage (see Figure 2.4) [159]. Anderson et al. [12,207] explain how learners achieve this high rate of performance improvement. They argue that people begin learning to carry out a task by dividing it into several sequential small parts; then, with experience gained over time, they merge those parts into a single action that can be carried out quickly. From the perspective of spatial knowledge development (Section 2.2.1), this stage is where people not only acquire landmark knowledge [74] but also start to form meaningful relations among the landmarks – they begin to acquire route and survey knowledge [209].

2.2.3.2 *Associative Stage*

In the associative stage of developing a skill, learners have already figured out how to complete a task and can concentrate on improving efficiency by incorporating slight changes to their reactions [213]. In contrast to the cognitive stage, learners in the associative stage primarily focus on how they are executing those actions [188,213]. As shown in Figure 2.4, another significant difference between these two stages is that the rate of performance improvement is lower in the associative stage [77,188].

Since learners have identified a set of actions to carry out a task effectively in the associative stage, people focus on rehearsing the set of actions learned in the cognitive stage that will eventually enable them to execute the actions more quickly and accurately [77]. In the context of spatial knowledge acquisition, in the associative stage, people begin to acquire survey knowledge [209] by forming associations among the landmarks learned in the cognitive stage. In other words, people start to perform like experts.

2.2.3.3 *Autonomous Stage*

The autonomous stage is the final stage of skill development (see Figure 2.4). In this stage, learners acquire a better ability to judge which stimuli are relevant, and the skills become automated; as a result, a lower degree of consciousness is required to select corresponding responses [77]. In this stage, the learners become experts who can perform tasks efficiently in terms of speed and accuracy [213], primarily because they automatically rely on the stimulus-response associations learned in the earlier stages of skill development. In the context of spatial knowledge acquisition, people develop survey knowledge of an environment that serves as a map-like representation of the environment consisting of landmark and route knowledge [74,209].

Whereas learners use declarative knowledge to perform tasks in the early stages, in the autonomous stage, learners develop procedural knowledge [12,123], which enables them to execute actions subconsciously and even in conjunction with other tasks [77,189]. Examples include riding a bike while having a conversation [213] or typing on a touch screen while dealing with other tasks [51]. People in the autonomous stage do not have to deliberately think about executing actions or selecting reactions to stimuli [189]; however, substantial time and practice are required to reach this point.

2.3 SUPPORTING SPATIAL LEARNING

Knowledge of the location of an item is often represented with reference to other objects or items. Several spatial reference systems provide inherent support to learners for learning spatial locations. This section presents commonly available techniques and strategies to aid spatial learning and recall.

2.3.1 Frames of Reference

Locations in real-world environments or GUIs are relative. It is difficult to describe the location of an item in a space without setting up some type of reference frame. People usually learn, organize, and communicate spatial knowledge (i.e., information about an object's location) by recognizing the spatial relations among the objects present in an environment.

Early research suggests that while learning the locations and spatial relations of objects in an area, people encode spatial information with reference to spatial frames [156,193]. Learning and remembering objects' locations is directly connected to learners' knowledge and interpretation of the surrounding area, which eventually supports establishing intrinsic reference frames [157,193]. The literature on spatial cognition usually divides the frames of reference into two categories that differ based on the perspective from which an object's location is viewed: egocentric and allocentric reference systems [134,156,166]. These frames of reference are presented in brief before discussing how they support people in developing spatial knowledge.

2.3.1.1 Egocentric

Egocentric frames of reference usually specify the location of an object or its orientation with respect to an observer that often includes "eye, head and body coordinates" ([153], p. 589). Since an egocentric reference frame represents an object's relation to the observer, it is also described as a viewer-dependent, relative, self-to-object, or first-person perspective-based reference mechanism. People use this reference system to perform navigation in daily life. Grech et al. described egocentric reference-based navigations as utilizing "direction (i.e., left-right) responses and actions independent of environmental cues, [where these] directional decisions are made at single or sequential choice points" ([85], p. 106) without depending on landmarks available in the environment. In other words, egocentric reference frames can help to develop route knowledge [195]. As shown in Figure 2.5: A, an egocentric reference frame is employed to specify routes from the start point 'a' to the destination point 'b' using a series of turns (e.g., turning left at the corner). Since the direction is specified from an observer's point of view, the same strategy would not work to reach the destination point if the starting point changed.

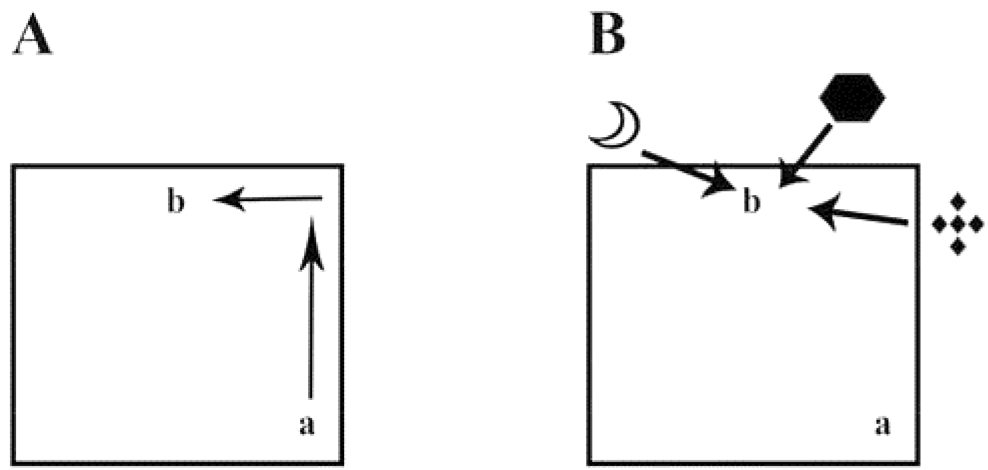

Figure 2.5: Frames of reference. A: egocentric, B: allocentric; adapted from ([85], p. 107).

Several prior works suggest that people rely on egocentric reference systems to cognitively represent and remember the locations of objects in a relatively small area, such as a standard-sized room [69,156,179]. For instance, Shelton et al. carried out a study where people learned two different views of a set of items in a room; later, they were asked to identify the items' relative directions from memory [192]. Speed and accuracy increased in the two familiar views compared to those displayed from new viewpoints, indicating that subjects developed spatial knowledge through the egocentric frame of reference.

2.3.1.2 Allocentric

Allocentric reference systems represent the location of an object with reference to other objects or reference frames such as landmarks available in an environment. Since this referencing mechanism utilizes external objects, it is also known as an object-to-object, third-person perspective, or intrinsic reference system [85,156]. As shown in Figure 2.5: B, the allocentric reference system relies on the combination of three landmarks or external identifiers to specify the location of point 'b.' Using those referencing points as cues, an observer can reach point 'b' even if the starting point 'a' changes location. Another example of an allocentric reference frame is the use of cardinal directions (e.g., east, west), as these cues remain unchanged regardless of an observer's orientation. One significant advantage of allocentric reference systems is the "flexibility of being able to locate novel points from various start locations as long as the external cues remain the same" ([85], p. 107).

Researchers have also extensively studied allocentric reference systems (e.g., [153,193]). Studies conducted by Shelton et al. [193] and Werner et al. [225] revealed that the arrangements or structures of objects present in an environment and their relative positions help an observer develop spatial memory of those objects. Therefore, allocentric reference systems enable people to develop survey knowledge [195] of an environment. The benefits of allocentric reference frames are apparent when we arrive at a decision point, where "we make judgments about the relative position of objects based on our memory of the location where they have previously been encountered" ([73], p. 2). For example, as Ekstrom et al. described, upon reaching a landmark in an area, "we could remember that our destination is positioned between this landmark and another one, sitting about 2/3 of the way from the 2nd landmark and at a 30° angle from the first one" ([73], p. 2).

Although these two referencing mechanisms are fundamentally different, research suggests that people often use them together [73,85]. McNamara et al. presented a novel framework stating that when developing knowledge of a new environment, people utilize their egocentric experiences along with the structural information of the new environment to identify an allocentric frame of reference to encode spatial information [153]. Mou et al. summarized the framework noting that "the spatial reference systems used in memory are anchored in the world, and in this sense are allocentric, even though they may be initially defined by egocentric experiences” ([156], p. 162). Therefore, this dissertation does not aim to differentiate between allocentric and egocentric reference frames; instead, it focuses on the general concept of reference frames, allowing people to develop knowledge of objects’ locations in computer interfaces.

2.3.2 Landmarks

In real-life, location learning benefits significantly from landmarks available in an environment [15,143]. Landmarks are easily identifiable and permanent objects in an environment that can provide useful reference frames for nearby objects [143]. Besides specifying and recalling objects’ locations in an environment [187], people can also use landmarks to determine their own location in the environment [118]. Tlauka et al. defined landmarks as “distinctive spatial features that, by virtue of their shape, colour, semantic value etc., have the potential to help individuals to orient and find the way around an environment” ([210], p. 305). The memory of landmarks is so prominent that it can help people navigate even when the landmarks are not visible [141].

The following sections briefly discuss the roles that landmarks play in developing spatial memory and the characteristics that define landmarks.

2.3.2.1 Functions of Landmarks

Research in cognitive psychology and urban planning shows the importance of landmarks in developing spatial memory [143,200]. Studies carried out on humans, animals, and even insects (e.g., [112,145,223]) indicated that landmarks act as reference frames to carry out spatial tasks, such as wayfinding and remembering locations. With a distance estimation task, Allen showed that participants performed better because of the presence of landmarks [9]. Siegel et al.’s spatial knowledge development framework is grounded in landmarks as they argued that people first acquire landmark knowledge upon arriving at and exploring a new area [195]. Golledge described two major roles that landmarks play in spatial tasks: organizational and navigational [83].

Organizational role of landmarks. Landmarks can provide general organizational information about an area by representing groups of objects or items together from that area [83]. Presson et al. [173] use Paris’s Eiffel Tower as an example: although it is just one landmark, it represents an abstract location of the whole city of Paris without revealing actual spatial locations of the places or objects in the city. Another example of landmarks' organizing role could be spatial reference points (e.g., familiar objects having cultural importance [181]) that allow people to recall nearby objects' locations in relation to those reference points.

Navigational support. Landmarks provide directional support allowing people to perform navigation in an environment [83]. The supports consist of "identifying choice points where navigational decisions are made, the origin and destination points, providing verification of route progress and influencing expectations, providing orientation cues for homing vectors and suggesting regional differentiating features" ([200], p. 40). While navigating along a path, landmarks provide memory cues that enable people to learn and remember locations in the path [8]. Sorrows et al. pointed out that "landmarks also enable one to encode spatial relations between objects and paths" ([200], p. 41). As a result, landmarks contribute to building a mental map of an area – the survey knowledge described by Siegel et al. [195]. People can use this knowledge to identify new navigation paths to reach a destination.

2.3.2.2 Attributes of Landmarks

Although landmarks need not be visual (they can also be cognitive), humans primarily rely on visual landmarks for wayfinding [137]; therefore, this dissertation limits the discussion to visual characteristics. Lynch [143], in his seminal work 'The Image of the City,' identified elements that people use to remember locations and navigate around cities. This section describes the elements Lynch identified and then summarizes other relevant attributes of landmarks found in the literature.

In 1960, Lynch carried out a study where he interviewed residents of three American cities (Los Angeles, Boston, and Jersey City) to understand how people develop spatial knowledge of their cities and what contributes to building that knowledge [143]. Lynch argued that the residents of a city develop a coherent 'mental image' [212] of the city by collecting and storing spatial information about objects they come across during their daily lives. He identified five elements that make an urban area memorable [143]. First, *paths* refer to the channels where a navigator can move - "streets, walkways, transit lines, canals, railroads" ([143], p. 41). Lynch also stated that "people observe the city while moving through it, and along these paths the other environmental elements are arranged and related" (p. 41). Second, *edges* are channels (e.g., a river) denoting the boundary of a district or a small area. Edges have organizing properties allowing people to "hold together generalized areas, as in the outline of a city by water or wall" ([143], p. 41). Third, *districts* are 2D areas traceable from the inside, such as a neighbourhood. Fourth, *nodes* are important points of interest in an area (e.g., a public building or a junction). Lastly, *point landmarks* are external elements or physical objects (e.g., statues or mountains) that can be seen from far away. Vinson [222] argued that all of these elements provide similar spatial benefits (e.g., orientation and navigation), so they can all be treated as landmarks.

Upon studying residents' cognitive images, Lynch [143] also described several characteristics of a landmark:

Singularity. Singularity makes the landmark visually distinct from other objects in an environment [143,200]. A building or a structure in an area can achieve singularity because of its "difference in size, shape, position, age, or cleanliness" ([200], p. 42) from surrounding structures. An example is the Thorvaldson building at the University of Saskatchewan. This iconic castle-like structure differs from other buildings located around it, so the Thorvaldson building acts as a landmark for nearby locations. Other unique visual structures, such as an intersection of roads, can also become landmarks due to singularity. Although landmarks have multiple features, Lynch argued that singularity plays the most important role [143].

Prominence. Another characteristic that turns an object or a structure into a landmark is prominence [143,200]. As Sorrows et al. [200] described, a prominent structure such as a building in a city environment can be a landmark that is "visible from many locations, or that stands significantly at a junction of roads" (p. 42). For example, the CN Tower in Toronto has this prominence feature: it is visible from many areas of the city, and it can provide spatial cues for city residents and visitors. Prominence is a visual characteristic, and researchers have divided landmarks into two categories based on their visibility [105,136,203]: global and local.

- *Global landmarks.* Landmarks that are always visible from all directions in an environment are treated as global landmarks. Since global landmarks are consistently visible, they can provide stable and useful reference frames for "traversing between two locations in separate occasions and along different routes" ([136], p. 90).

- *Local landmarks*. Local landmarks are only visible "at limited locations and from restricted perspectives" ([136], p. 90). Research has indicated that local landmarks can improve route choice accuracy during a wayfinding task by providing location information [114].

Both global and local landmarks can provide spatial benefits for navigation and remembering locations [136]. A study carried out by Steck et al. [203] revealed that route choice accuracy was significantly hampered when either global or local landmarks were absent. Researchers also looked at the differences between global and local landmarks; due to the differences in visibility, global landmarks primarily provide orientation information while local landmarks typically provide locational knowledge of objects [105,136].

In addition to these attributes, several other features can define an element in an environment as a landmark. For example, the meaning that an object portrays, the way it is typically used, or even its subjective importance [16] are features that can, together or in isolation, make an object a landmark. Whatever attributes landmarks have, they serve a common purpose: they provide frames of reference to learn and remember the locations of nearby objects.

2.4 SUPPORTING SPATIAL LEARNING IN GUIS

As described in earlier sections, spatial memory is a cognitive ability that allows people to learn and recall the locations of objects and places in daily life [121,170]. Researchers have carried out studies exploiting people's spatial memory to facilitate location learning and revisitation in various computer interfaces [72,177,187]. The following sections briefly summarize techniques and measures HCI researchers have used to support users' spatial learning and recall in GUIs.

2.4.1 Spatially Stable Layouts

The stability of a GUI's layout is an important precondition to developing spatial memory. GUIs dynamically change the positions of graphical elements, such as scrolling interfaces or windows that "re-flow" icons when the size of the window changes, which can hinder the natural process of developing spatial memory of those elements [50]. As a solution, researchers have investigated spatially stable representations that display all elements at once (similar to Figure 2.6) to benefit spatial memory development [49,183]. Previous work has looked at two types of interfaces: scrolling and hierarchical.

2.4.1.1 Spatial Layouts vs Scrolling Interfaces

Several research projects in HCI have explored spatially stable layouts as a way to improve spatial learning and revisitation performance, primarily in linear documents that usually rely on scrolling-based interactions. While investigating people's reading behaviours in papers and digital documents, O'Hara et al. [164] saw that all participants separated pages from the provided page bundle and placed them on a table to get an idea of the document's overall structure. They also found that scrolling techniques (standard for digital document navigation) were "irritatingly slow and distracting" ([164], p. 338) in navigating online documents.

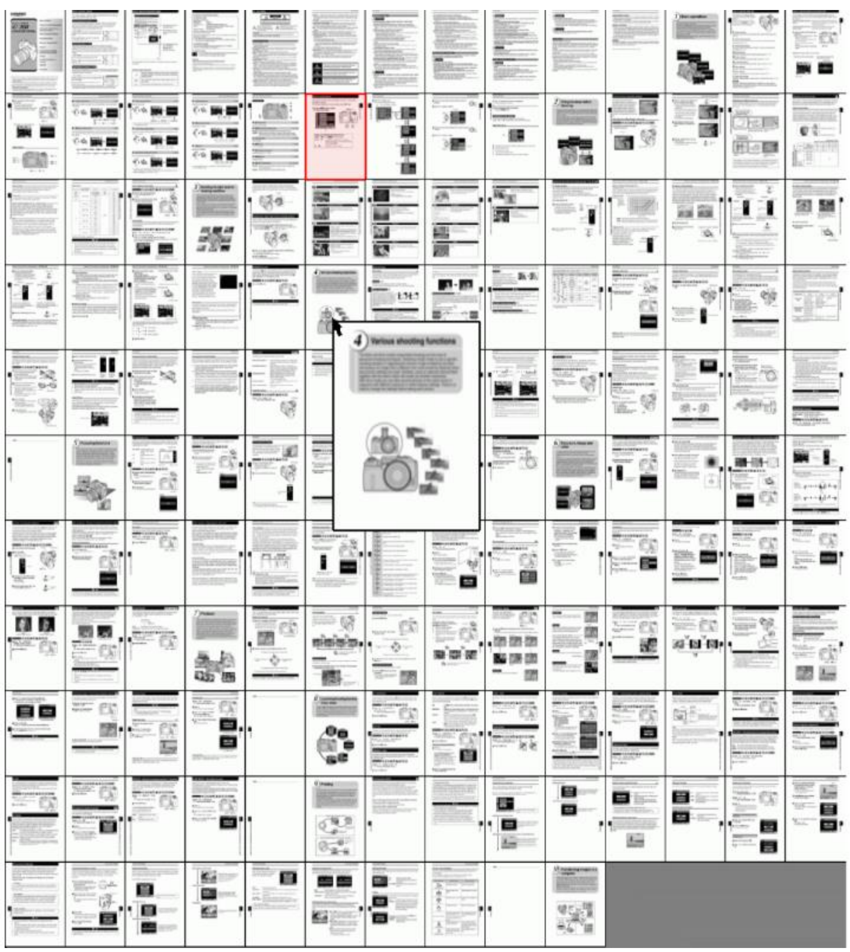

Figure 2.6: Space-Filling Thumbnails system showing pages as thumbnails of an entire document on a screen [49].

Cockburn et al. [49] developed Space-Filling Thumbnails (SFT) in order to improve digital document navigation performance and facilitate spatial learning. Instead of using traditional scrolling, SFT employs a spatially stable overview that displays all pages of a document as thumbnails arranged in a grid. As shown in Figure 2.6, each page of a document appears at a fixed location on the screen, which allows users to build spatial memory of the document. Results indicated that SFT outperformed six other document navigation techniques [49], particularly when revisiting previous locations. In a follow-up study, Gutwin et al. [93] confirmed the value of STF's spatially consistent overviews for document navigation in non-laboratory settings.

Figure 2.7: A ListMap system displaying fonts in a spatially stable grid layout [92].

Another example of a spatially stable interface is Gutwin and Cockburn's ListMaps [92]. They transformed a linear font menu into a 2D grid where each cell represented one font name from the list (see Figure 2.7). ListMaps significantly improved item revisitation performance compared to a traditional scrolling Listbox since ListMaps users could better use spatial memory with the stable layout. However, novice users found ListMaps more difficult for initial search [92], possibly because of the small size of the grid cells or the shortened font names (only a portion of the full name was visible by default; the full name was only displayed upon mouse-hover).

2.4.1.2 *Spatial Layouts vs Hierarchical Menus*

Many GUIs organize commands into hierarchical structures such as menus, ribbons, or tabs. Hierarchical structures provide benefits, such as a categorical organization, that help novices find commands easily [104,182]. However, Cockburn and Gutwin suggested that "broad and shallow" arrangements of commands could benefit users' spatial learning of commands [48]. As a result, researchers began to explore alternative menu designs with the aim to leverage spatial memory and support spatial learning.

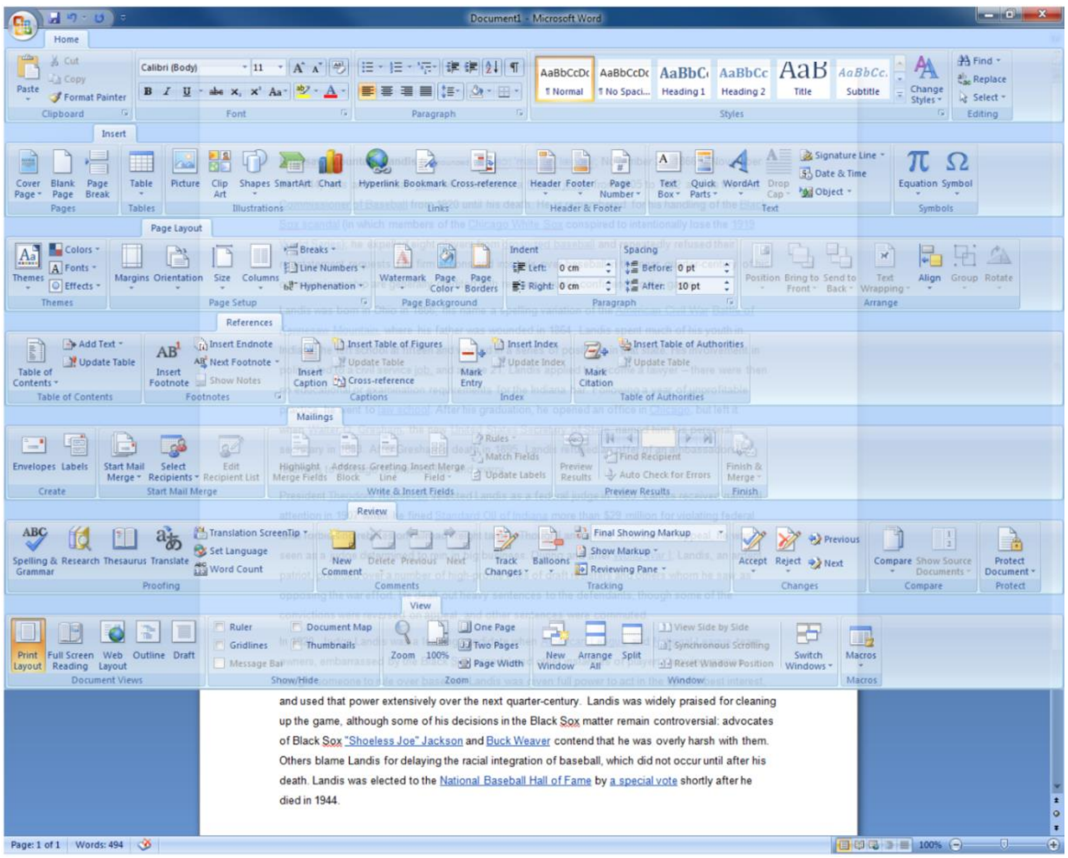

Figure 2.8: CommandMap interface showing all commands of MS Word's ribbon [183].

One approach attempts to support users' spatial learning by completely flattening the traditional command hierarchy to display all commands simultaneously, similar to the techniques described in Section 2.4.1.1. As shown in Figure 2.8, Scarr et al. [183] developed the CommandMap interface that uses a full-screen overlay to display all the available commands of Microsoft Word. Since each of the commands appears at a specific location on the interface, users only need to learn the location of a command (e.g., the “copy” command is at the top left corner) instead of learning a multi-level ribbon structure.

Scarr et al. [183] compared CommandMap’s command selection performance with traditional menus and ribbons. Although results did not show differences for novices, experts selected commands significantly faster than standard menus. Besides leveraging spatial memory, CommandMaps [183] enabled rapid command selection by reducing the total number of steps required to select a command. Similar results were observed when CommandMaps were tested in realistic tasks with two applications: Microsoft Word and Pinta (an image editing program) [184].

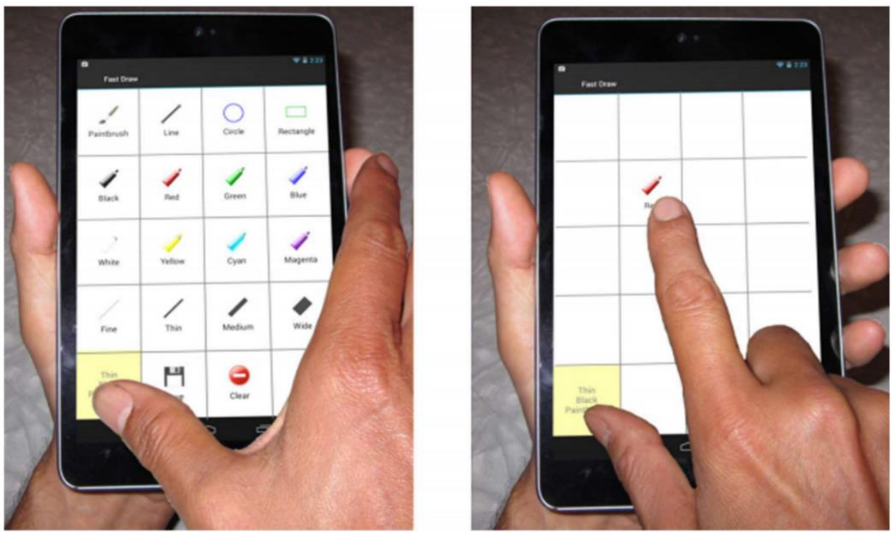

Figure 2.9: FastTap selection: (left) visual search by a novice, (right) rapid selection from memory by an expert without waiting for commands to appear [95].

The performance benefits of flat all-items-at-once menu designs in desktops have inspired researchers to investigate similar menu designs on other platforms. For example, Gutwin et al.’s FastTap [95] uses a spatially stable grid interface to provide rapid command selection for expert users of touch tablets (see Figure 2.9). Novice users of FastTap can learn the desired command location via visual inspection after invoking the menu with a thumb. Once learned, expert users can use a chorded thumb-and-index-finger touch (Figure 2.9) to quickly select a command by recalling its location from memory [95].

2.4.2 Augmented GUIs with Visit Marks

Several researchers have attempted to improve spatial revisitation in GUIs by augmenting the widgets present in an interface with visual cues about past locations, such as highlighting locations that people previously visited [7]. One early technique – Hill et al.'s 'read wear' – showed user visits as a histogram in the scrollbar of a linear document reader [100]. Other techniques, such as Mural Bars [152] and AlphaSliders [3], also displayed document contents as miniature visualizations or alphanumeric characters in the scrollbar area to provide users with a better spatial understanding of the corresponding document. Inspired by these scrollbar augmentation methods, Alexander et al. [7] proposed the 'Footprints Scrollbar' that displays temporary marks to indicate recently- and frequently-visited document locations in the scrollbar. As shown in Figure 2.10, users can quickly revisit locations they visited previously by clicking on the corresponding colour marks. Results also showed that these visual cues could act as landmarks, helping to decrease revisitation time [7].

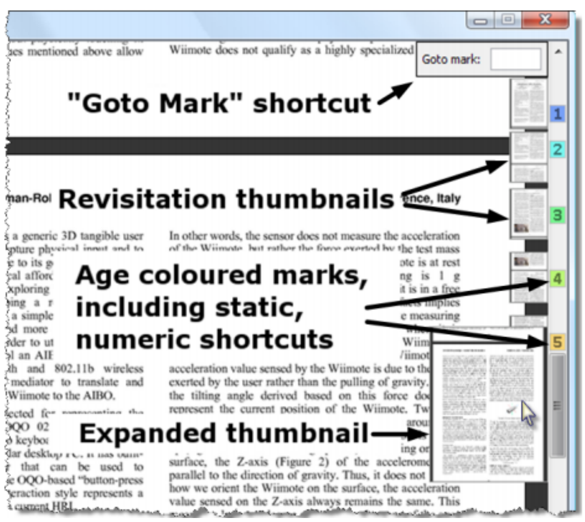

Figure 2.10: Footprint Scrollbar – colour markers to enable revisitation [7].

Media players are another type of interface where augmentation is commonly used, although previous work does not explicitly focus on revisitation. To support exploration within a video, researchers augmented the timeline slider of the media player to show visual highlights to represent personal [5] or crowd [124,231] navigation history. Chen et al.'s [47] Emo Player annotates the media player's slider with different colours based on the characters' emotional states. Schoeffmann et al.’s video explorer [190] augments the slider with the dominant colours of each frame to help users navigate and explore video content.

However, a problem with techniques based on these visualizations is that there are situations where a user may want to revisit (or visit) a location not shown in the widget. For example, suppose a user watches an entire video clip. In that case, all locations are visited equally during the initial playback, which makes it harder to return to a particular scene. Similarly, many visited locations may not appear in a Footprints-style augmentation. For example, some locations have not been seen often enough (or for long enough at each visit) to appear in the visualization, while some locations that go unvisited for a time can disappear from the list of recent items.

These problems arise because the techniques mentioned above depend on the system to remember and visualize important document locations. A different approach is to rely on the users to remember these important locations, which is possible if they are given proper resources (e.g., landmarks) to exploit their spatial memory abilities.

2.4.3 Landmarks in GUIs

Landmarks play a crucial role in supporting location learning and recall (see Section 2.3.2). They can also aid spatial memory development of GUIs by helping users to mentally consolidate the data they see, resulting in an overall spatial understanding. Research in HCI that has explored GUI objects as landmarks can be divided into two groups: natural landmarks and artificial landmarks.

2.4.3.1 Natural Landmarks

In order to make graphical interfaces more expressive and aesthetic, designers often augment GUIs with various graphical features such as colours, symbols, images, or icons. Although these graphical features are not explicitly identified as landmarks in design guidelines [162,194], they can act as reference frames for spatial memory development.

There are several naturally occurring landmarks on devices that researchers have exploited on different occasions. For example, Gutwin et al.’s FastTap [95] (see Figure 2.9) uses the corners and bezels of a tablet device as landmarks to create a strong frame of reference for the commands displayed on the device. Schramm et al.’s Hidden Toolbars [191] used the edges and corners of touch tablets as landmarks to organize four toolbars. Although these landmarks can provide reliable reference frames for the commands placed near them, they are less likely to provide the same level of spatial benefit for large devices since many commands will be located far away from these landmarks.

In areas where these prominent natural landmarks are absent (such as the middle of a large touch screen), other natural objects such as a user's own hands can be utilized as landmarks. Uddin et al. [220] demonstrated two bimanual command selection techniques called HandMark Menus for large multi-touch tables. As shown in Figure 2.11: left, one version of the technique places command icons in the spaces between a user's spread-out fingers. This technique uses the hand as a clear external reference frame – once the locations of different items are learned, people can use the proprioceptive knowledge of their hands as a frame for setting up the selection action even before their fingers are placed on the touchscreen [220]. This technique can be used with both hands to increase the number of available items. A second version (Figure 2.11: right) accommodates larger command sets by placing blocks of commands between the thumb and first finger, with different sets accessible through different finger combinations.

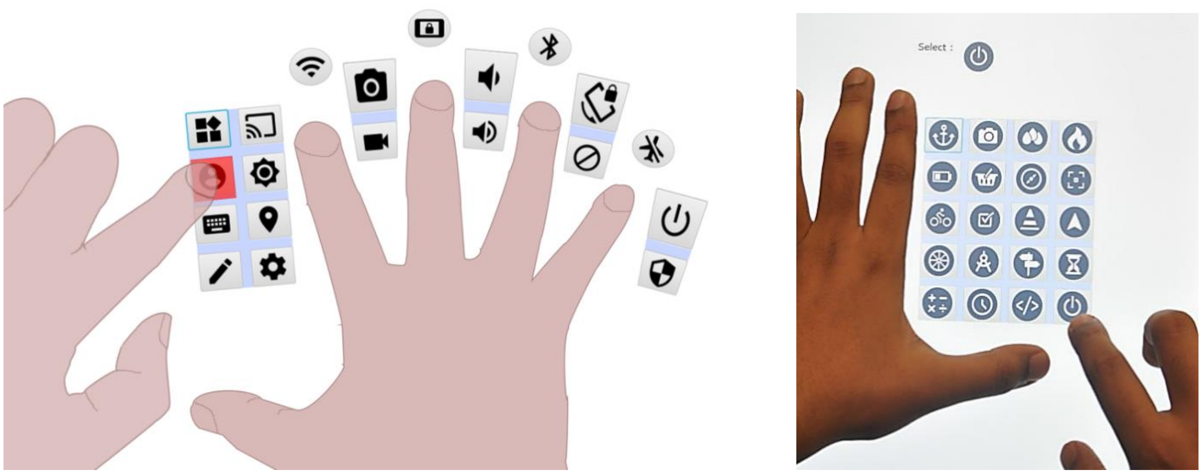

Figure 2.11: HandMark Menus. Left: selecting a command from the HandMark-Finger menu; right: selecting a command from the HandMark-Multi menu [215].

In addition to large touch tables, the use of hands as a potential reference frame has been explored in 3D virtual environments. For example, Hinckley et al. [101,102] demonstrated that people performed better in 3D object manipulation tasks because they could successfully coordinate between two hands since one hand provided a reference frame for the other hand to manipulate objects (e.g., aligning them). Gustafson et al.'s [91] Imaginary UIs also enabled users to perform touch-less gestures with one hand inside an imaginary frame created by using the other hand as a landmark. However, these landmarks are not useful in traditional desktop interfaces where indirect pointing devices (e.g., mice) are predominantly used.

2.4.3.2 Artificial Landmarks

When natural landmarks are insufficient, artificially created visual elements can serve as potential landmarks [17,200]. Artificial landmarks such as colour [7] can help people quickly revisit an intended location, and shape has been used to give an object a memorable "visual ID" [135]. For example, Alexander et al. applied this idea in their Footprints Scrollbar (see Figure 2.10); this system places coloured marks in the scrollbar area for recently visited locations [7]. An evaluation showed that these visual cues could act as landmarks, decreasing navigation time when revisiting locations. Other techniques that have also successfully applied artificial landmarks in the form of visit marks include Skopik and Gutwin's "visit-wear" [198], City Lights [232], Halo [31], Wedge [90], Visual Popout UIs [78], and the Canyon visualization [106]. These visit marks are shown to aid in revisiting previously seen items within the interface.

Schramm et al.'s [191] Hidden Toolbar (Figure 2.12) shows colour markers in a thin line along the four edges of a tablet; each colour represents a particular icon or command. These marks act as landmarks and aid users perform rapid selections by swiping outward on those marked locations.

Figure 2.12: Hidden Toolbar interface – using coloured border regions as landmarks to help users remember locations [191].

Researchers have also augmented media player interfaces using representations of navigational activities as landmarks to aid video exploration. For example, SceneSkim [167] adds reference points for different video locations, and Video Digests [168] represents sections/chapters with navigable markers. Some of these techniques extract information from the content and present it as colour distributions [47,190] in the slider to support exploration and revisitation.

The techniques described above often provide performance benefits compared to standard and un-augmented interfaces; however, none explicitly identified the graphical features as landmarks and explored their effects on spatial memory development and revisitation. Therefore, in this dissertation, I concentrate on the effect of landmarks in developing spatial memory of GUIs.

CHAPTER 3

SPATIAL MEMORY DEVELOPMENT IN STANDARD GRAPHICAL INTERFACES[6]

Citation: Md. Sami Uddin, and Carl Gutwin. 2021. The Image of the Interface: How People Use Landmarks to Develop Spatial Memory of Commands in Graphical Interfaces. In *Proceedings of the 2021 Conference on Human Factors in Computing Systems,* ACM, New York, NY, USA, 17 Pages.

Contributions and achievement: Under the supervision of Dr. Carl Gutwin, I designed and conducted the study, performed data analyses, and prepared the manuscript with study findings. This publication has received an ***Honorable Mention Award*** at CHI 2021!

User interfaces visually represent tools, commands, or control mechanisms by placing them in certain locations of an interface. Although an interface presents a large number of visual commands, often arranged in complex structures such as hierarchical menus or toolbars, experienced users of an interface (after adequate practice) can easily access them by quickly recalling their locations from memory. However, very little is known about how such spatial memory of commands is developed in standard computer interfaces. Since landmarks present in real-life environments substantially influence spatial learning, this chapter aims to uncover if landmarks are present in standard graphical interfaces, and whether they contribute to users' location learning and revisitation performance in these interfaces. I carried out a study

[6] The main content of Chapters 3, 4, and 5 contain published papers that have not been edited. However, at the start and end of each chapter, I include short sections to establish the problem, motivation, and contributions of each manuscript in terms of the overall context of the dissertation. In Chapter 3, this context is provided in introductory sections 3.1 and 3.2, and concluding section 3.4; the complete manuscript for the published paper appears in section 3.3.

investigating interfaces of four commercially available applications; results indicated that existing features and design elements in GUIs (e.g., an interface's layout, corners) do act as landmarks, and that users relied on landmarks to learn and recall the locations of commands.

3.1 PROBLEM AND MOTIVATION

Graphical interfaces provide an easy-to-use interaction facility for users by making widgets visually available on a screen. Users, particularly novices, who are new to a computer interface, must perform their desired tasks in an application by visually searching for and then interacting with the available menu-based commands. Although most computer applications show a large number of commands in their interfaces (e.g., Microsoft Word and Adobe Photoshop have several hundred), frequent users of those GUIs can find commands quite efficiently to perform tasks. These expert users can even find commands without carrying out any visual search, indicating that users develop spatial memory of those commands. It is, however, not clear how memorable these interfaces are and what contributes to users developing spatial memory of commands in GUIs.

Location learning in real-life greatly benefits from available landmarks in an area. Landmarks provide a spatial reference frame for people to remember places and locations around them, allowing them to conveniently revisit locations and navigate an area [143]. Researchers have developed several prototype interfaces [81,95,132] leveraging the corners and bezels of the small screen as landmarks to support spatial learning. For large touch screens (e.g., tabletop), users can even use their hands as landmarks [216] to learn and remember a large number of commands. However, such landmarks are uncommon in currently available standard desktop GUIs. Existing features in graphical interfaces, such as an interface's layout, edges or corners, may provide benefits while learning locations in GUIs, but design guidelines [162,194] do not explicitly discuss the use of any landmark in GUIs. Although it requires practice, people eventually become fluent in performing memory-based actions (e.g., command selection or location revisitation) in these interfaces. Therefore, it is important to know if landmarks are present in these interfaces to support the development of spatial memory of commands so that interface designers make graphical interfaces' conceptual visual models more vivid and quickly memorable.

The problem addressed in this chapter is that there is a clear gap in our understanding of how spatial memory development takes place in existing graphical interfaces and whether landmarks are available in these GUIs that help users learn and recall the locations of commands.

3.2 SOLUTION AND STEPS TO SOLUTION

To explore these issues, I designed and carried out a study to understand the memorability of different computer interfaces and the ways users remember contents in the visual representation of an interface. The study primarily focused on discovering how people learn and remember various commands in standard interfaces and if they employ any visual element or feature from the interfaces as landmarks in that spatial learning process.

Most of the tasks people perform on computers primarily involve interacting with graphical interfaces of various applications. For example, editing a text document requires the MS Word interface; reading a PDF document involves the Adobe Acrobat Reader, and editing a photo requires the Photoshop interface. Though these graphical interfaces present commands visually on the interface, each of these interfaces follows a unique layout or structure to represent the commands. While performing tasks in an interface, people navigate through and interact with the interface's elements, which eventually enables people to develop a mental image [212] of that interface. Analyzing those mental images will reveal what contributes to the development of spatial memory of the commands in GUIs and how memorable those interfaces are.

3.2.1 Research Questions

This work investigated the mental images people have in their minds about different interfaces to understand how landmarks benefit spatial memory development in standard computer interfaces. I designed a semi-structured interview-based study to answer the following three questions:

- Do users develop images of an interface: that is, do they develop spatial memory of commands and interface features and remember their locations?

- How strong and accurate are the images of an interface: that is, how accurately can users remember the locations of commands?
- Do people use landmarks to remember the locations of commands and navigate in an interface, and what are those landmarks?

3.3 MANUSCRIPT A

Graphical User Interfaces present commands at particular locations, arranged in menus, toolbars, and ribbons. One hallmark of expertise with a GUI is that experts know the locations of commonly-used commands, such that they can find them quickly and without searching. Although GUIs have been studied for many years, however, there is still little known about how this spatial location memory develops, or how designers can make interfaces more memorable. One of the main ways that people remember locations in the real world is landmarks – so we carried out a study to investigate how users remember commands and navigate in four common applications (Word, Facebook, Reader, and Photoshop). Our study revealed that people strongly rely on landmarks that are readily available in the interface (e.g., layout, corners, and edges) to orient themselves and remember commands. We provide new evidence that landmarks can aid spatial memory and expertise development with an interface, and guidelines for designers to improve the memorability of future GUIs.

3.3.1 Introduction

Graphical interfaces provide a visual two-dimensional representation of commands arranged in menus, toolbars, and ribbons, where each command appears at a specific location in the interface [130,183]. In order to locate commands, novice users of those GUIs must perform a slow visual search among the icon-based commands [187], primarily because they have not memorized their locations. In contrast, frequent users of an application can easily find commands from memory without relying on visual search, even though many commercial applications such as Adobe's Photoshop or Microsoft's Office suite have hundreds of commands in their interfaces [50]. For example, a frequent MS Word user knows that 'Paste' is located near the top-left corner of the interface, and that it is the leftmost item in the Home ribbon. Spatial memory is a powerful way to enable expert performance with GUIs [54,92,184]. One benefit of human memory with spatial locations is that people can quickly and accurately retrieve the locations of frequently-visited items from memory [170]. Even with a blank ribbon, Scarr et al. [183] showed that people could accurately recall ~30 commands (50% of the commands they used in MS Word) – indicating that users had developed strong spatial memory of those commands [187]. However, we do not precisely know how this spatial location memory is formed in GUIs.

One mechanism that helps people learn spatial locations in real life is the landmarks in an environment [61,170,195]. Landmarks help people to remember the places and locations they have visited earlier [143]. They also play a critical role in helping people move towards "survey knowledge" [195,209] – that is, a mental map of an area (also called a *cognitive image* [103,143,212]). Landmarks provide spatial anchors in the mental map that help people remember object locations. Upon studying the cognitive images of three American cities back in 1960, Lynch [143] suggested that the citizens of a city develop a "mental image" of the city because of their spatial knowledge. He also identified five features of cities that can be considered landmarks [143,222]: paths (obvious routes between places), edges (physical boundaries such as rivers), districts (obvious regions within the city), nodes (strategic points of interest such as intersections), and point landmarks (specific objects such as natural features, buildings, or monuments).

Landmarks are also present in GUI systems and can aid spatial learning [187]. Researchers have designed several GUIs [81,95,132,191] that leverage the outer corners and bezels of small devices (e.g., tablets) as landmarks. Even a user's own hands and fingers [91,101,102,220,228] have been used as landmarks to benefit spatial learning. However, there are only a few natural locations that can be used as landmarks in GUIs; in the absence of natural landmarks, artificially-created landmarks can also improve the performance of spatial tasks [80,154,205,218,219]. Artificial landmarking features are uncommon in GUIs, perhaps because of the risk of distraction [175,198]. Still, frequent users are fluent in performing memory-based actions [183] – indicating that, over time, they have developed spatial images (i.e., mental maps) of the GUIs. It would be useful for designers to understand how this kind of expertise works – to know what landmarks are present in these GUIs that help develop spatial memory, and what designers can do to make GUIs more easily memorable.

In order to identify the role that landmarks play in building spatial memory of commands in current GUIs, we explored the mental images that frequent users develop of the interfaces they use. Lynch [143] successfully demonstrated a method to elicit the cognitive images of a city from its inhabitants; to determine the images of GUIs that are in expert users' minds, we carried out a study adapting Lynch's [143] method to explore four conventional GUIs: Microsoft Word, Facebook, Adobe Acrobat Reader, and Adobe Photoshop.

Our study showed that frequent users of the four applications do have strong images of the GUIs in their minds, and they could easily recall the locations of commands using several types of landmarks. We found that, among the four GUIs, Word and Facebook were relatively more memorable (over 76% accuracy in verbal descriptions) than Reader and Photoshop (below 44% accuracy). Accuracy was even higher when participants carried out location pointing tasks on a sketch or a "washed-out" version of the UI (73% accuracy overall). We also found that people relied heavily on landmarks (e.g., GUI layout, command groups, internal and external corners, and icon visuals) to orient themselves to the GUIs and remember command locations – indicating that landmarks can aid in developing spatial memory of commands in GUIs.

Our work provides three contributions. First, to the best of our knowledge, we are first to show that frequent users of GUIs can develop vivid spatial images of GUIs, and we provide new evidence that landmarks readily available in GUIs can aid spatial memory and expertise development. Second, we identify four types of landmarks in GUIs that users can rely on to remember command locations correctly. Last, we provide guidelines for designers who want to make future GUIs more easily memorable.

3.3.2 Background and Related Work

In the following sections, we review the psychological aspects of spatial memory and landmarks, then report how spatial memory and landmarks have been leveraged in GUIs.

3.3.2.1 Spatial Memory, Landmarks, and Cognitive Images

Spatial memory is a crucial human cognitive ability that allows learning and recalling the locations of objects and places in daily life [121,170]. Human memory has long been studied in psychology [23,25,46,61,66], where researchers have looked at how spatial memory is developed [170,208] and suggested that spatial learning in the real world greatly benefits from the landmarks in the environment [143,195]. Landmarks are readily identifiable and stable features or objects in a space that are easily separable from their surroundings and that can provide a frame of reference for other locations [143]. Landmarks can be both natural and human-made: for example, in a city area, a park or a prominent building can be a landmark, allowing people to remember nearby locations. Landmarks can be divided into two categories based on their visibility: global and local [136,203]. Global landmarks mainly provide orientation knowledge as they are visible from almost all regions of an environment, while due to limited visibility, local landmarks aid object locations recall in a specific area [105,136].

Siegel et al. [195] suggested a model that described spatial memory as a combination of landmark, route, and survey knowledge. After entering a new area, people naturally start learning the locations of objects with reference to other prominent objects (i.e., landmarks) [74,84,97] – which forms landmark knowledge [74]. Once people are familiar with the area, they begin to navigate between known landmarks and obtain route knowledge [209]. With further experience, people attain survey knowledge, where they have a complete understanding of the area along with its landmarks. Survey knowledge serves as a mental mapping of the items present in an environment – i.e., a cognitive image [103,212]. This cognitive image provides people with the spatial information required to recall an item or perform navigation in that area.

Lynch [143], in his seminal work 'The Image of the City' investigated the cognitive images of three cities. He argued that the dwellers of a city collect and store distinct objects they come across during their daily lives in their memory, forming a coherent mental image of the city. He identified five categories of objects that make an urban area memorable [143]: *paths*, channels where a navigator can move; *edges*, channels denoting the boundary of a district; *districts*, 2D areas traceable from inside; *nodes*, crucial points of interest; and *point landmarks*, external elements that can be seen from a distance. In this paper, we follow Lynch's [143] methods to explore the cognitive images of four GUIs and determine what landmarks are present in GUIs and how they aid spatial users' memory development.

3.3.2.2 Leveraging Spatial Memory in GUIs

Spatial memory has long been exploited in GUIs to enable memory-based user actions [72,177,187] and to support users' transition from novice to expert [51,127,235]. For example, Scarr et al.'s CommandMap [183,185] showed that spatially constant command placement in desktop UIs facilitated better command retrieval rate, even for real-world tasks [184]. The advantage arose because users could leverage spatial knowledge to recall the locations of commands from memory [52,183,218]. Other researchers have similarly shown that spatially-stable organizations that flatten command hierarchies to show all commands at once can improve revisitation efficiency compared to linear lists and menus [92] [49]. Research has also indicated that learning and recall performance can increase if command interfaces use meaningful icons [45].

Spatial memory can benefit touch and multi-touch interactions as well. For example, spatially stable command structures (i.e., each item in a cell) can improve selection speed in tablets [81,94,95,217], in smartwatches [132], in smartphones [233,234], and digital tables [220]. Even in large environments, spatial memory has been found to be useful (e.g., in VR [79], with wall displays [115], and large tables [116]). However, one basic question remains unanswered – what contributes to spatial memory development in standard GUIs that people commonly use.

3.3.2.3 Landmarks in Graphical Interfaces

There are two main categories of landmarks that researchers have exploited to aid spatial memory: natural and artificial. Natural landmarks in a GUI setting (e.g., screen corners) can offer support for spatial development [95,218], and researchers have already leveraged the corners and bezels of small touch devices to place commands [95,132,191]. In larger settings, however, these real landmarks become weaker, as locations are often far from the landmarks. In such cases, natural elements such as the users' own hands and fingers can offer spatial support. For example, systems by Hinckley et al. [101,102], Gustafson et al.'s [91] Imaginary UIs, and Uddin et al.'s [217,220] HandMark Menus use the non-dominant hand as a reference frame [228] for another hand to perform tasks efficiently. Yan et al. [229] showed that even the user's body could be a landmark.

In the absence of natural landmarks, digitally created visual objects can act as landmarks [218]. Vinson [222] and Sorrow et al. [200] suggested guidelines to design such landmarks, particularly for large environments. Mou et al. [157] leveraged the inter-object tie as a landmark in tabletops. Researchers used colour marks [7,198], and random icons [219] in the scrollbar of a document reader to show that landmarks could improve within-document revisitation, even in very long documents [154]. Also, artificial landmarks have been found to be useful in VR [79,80] and touch typing [205]. However, research suggested that overuse of landmarks might be distracting and can hamper performance [175,198].

Surprisingly, standard GUIs do not explicitly use artificial landmarks, but even so, experts with a GUI can locate commands accurately from memory [183]. Therefore, we carried out a study to explore what landmarks exist in current GUIs.

3.3.3 Study: The Image of the Interface

To understand the role of landmarks in developing spatial memory of commands in regular GUIs, we carried out an interview study with users of four interfaces. The following sections present the GUIs used and methods followed in the study.

3.3.3.1 Study Interfaces

The study concerned learning and recalling locations of commands in standard GUIs. We chose the interfaces of four popular desktop applications for our study: Microsoft Word, Facebook, Adobe Acrobat Reader, and Adobe Photoshop. We chose these four not only because many people use them frequently, but also because they provide variation in terms of the type of tasks carried out with the UI (e.g., photo editing, document processing, and social media), the number of commands (from a few to a few hundred), and the primary command arrangement (e.g., ribbons, menus, and toolbars). The versions of the applications that we studied are those from August 2018.

Word [237], a popular and well-known document processing application, arranges hundreds of commands in tab-based ribbon menus that usually appear at the top of the GUI. A large number of commands and the multiplexed ribbon menus create a unique challenge to remember command locations and navigate in the interface. *Facebook* [238], a popular social media platform, is a representative of complex web-based applications. Although it seems to have relatively few commands, Facebook utilizes the whole window to arrange its commands, which may make it more difficult to recall command locations. *Reader* [239], a well-known document viewing application, is the simplest interface among the four in terms of command number and arrangement. We were interested to see how people form location memory in a simple interface. Last, *Photoshop* [240], a popular photo editing and design application, uses a slightly different layout than the other systems – a canvas for editing photos in the middle and commands around the left, top, and right of the UI. A large number of commands are visible in the UI, possibly making it difficult for users to learn and remember their locations.

3.3.3.2 Study Method

Previous literature suggests that expert users of an interface can recall command locations from memory (without visual search) [95,183] – indicating they may have developed a *cognitive mapping* [46] of those commands. As a result, they can easily visualize an image of the interface in their mind. We therefore planned to elicit and analyze users' interface-images to explore what helps people remember commands and navigate in GUIs. Kevin Lynch, in his seminal work 'The Image of the City' [143], successfully elicited images of a city from its inhabitants' minds and revealed how people develop spatial memory of a city. Even after 60 years, architects and urban planners rely on this method to evaluate the visual perception of urban spaces. Therefore, to explore users' mental images of GUIs, we designed an interview study that adapted the methods used by Lynch.

In preparation for the study, we conducted an initial inspection of the four interfaces to understand their layouts. The aim was to check the structures and the arrangements of the available commands and also to identify their appearance, perceptibility, and the available landmarks in the GUIs so that later we could validate participants' responses. The main study consisted of semi-structured interviews with a small sample of people who use the four interfaces frequently. We aimed to elicit the images of the GUIs from users' memory, so the interview involved users describing the interface layout, the available commands, and the locations of commands they frequently use, along with nearby commands. It also included tasks such as drawing sketches of the GUIs, and identifying command locations on the sketches and on washed-out images (see Figure 3.1). The study ended with users carrying out walkthroughs (i.e., step by step instructions) to locate commands.

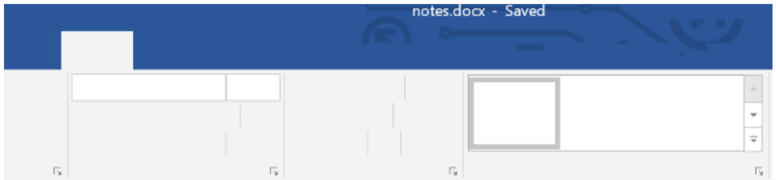

Figure 3.1: Part of the washed-out image of the Home ribbon of Word interface used in the study.

3.3.3.3 Interview Questionnaire and Tasks

Our semi-structured interview included the following questions (adapted from Lynch [143]) and tasks:

1. What image comes to your mind first when you think of or visualize the interface of ___?
2. Describe the interface of ___. What are the commands or tools available, and where are they located in the interface? [We asked for details when needed.]
3. Draw a sketch of the ___ interface, as if you are describing it to a person who has never seen it before. Try to cover the main features of the interface. [We took notes of the drawing sequence.]
4. Name 5 commands from this application that you frequently use or that you find most distinctive. They may be small or large, but should be commands that you can easily identify and remember.
5. (The following questions are repeated for each of the 5 commands that the participant provided)
   a. Describe the location of __ command in the interface. Picture yourself locating it and describe the path you would take and the items you would see to locate it.
   b. How did you remember the location of the command? Did you use any landmarks for it? Also, if you were uncertain about the command's exact location, what strategy would you use to locate it?
   c. Name the commands around <the chosen command>. You can describe the colour or shape of these items if you cannot recall exact names.
6. (Location pointing task: also repeated for the 5 commands that the participant provided)

    a. Show the command's location in your sketch.

    b. Point to the command's location in the washed-out image.

7. Walkthroughs for four imaginary tasks (Table 3.1) using real snapshots of the interfaces.

Table 3.1: Tasks used in walkthrough generation.

Word	**Facebook**	**Reader**	**Photoshop**
From Home ribbon, select insert an equation	From Home page, open friend list of a friend	Highlight texts in the document	Select brush and yellow colour
Change line spacing	Post a photo to own timeline	Add signature to the document	Turn on the 2nd layer
Insert a table	Delete that post	Fit a full page to the window	Add text in blue colour
Change the zoom level of the text.	Write a comment to a post in the News Feed	Go to a specific section of the document	Select all items then use 'left alignment'

As in Figure 3.1, washed-out images of the GUIs were created by removing the commands from a screen snapshot, keeping the outline and boundaries intact. This was done to see whether people use structural features as landmarks to locate commands. For the walkthrough tasks, we showed participants actual snapshots (printed on A4 paper) of the GUI. For each system, we came up with four tasks (see Table 3.1) and asked the users to generate step-by-step instructions to locate the required commands to perform those tasks. In the Word interface, for example, one task was *"Insert a table."* We analyzed the answers to look for any mention of landmarks, spatial locations, or frames of reference, in order to determine what landmarks were used for locating commands and navigating within the interface.

3.3.3.4 Participants and Apparatus

We recruited 20 people (5 women), 5 for each interface, ages 18-43 (mean 27.2) from a local university. Most of the participants were students (18), and two were working professionals. All were self-reported frequent users of the respective GUIs (daily or almost daily usage: 16; several times a week: 3), except for one who was only somewhat familiar with the Photoshop interface and used it about once a week. Table 3.2 summarizes their experience per interface: sixteen participants had been using these interfaces for over three years and one had less than one year of experience. Participants used these GUIs either on laptop or desktop computers (only one used Facebook on a smartphone). The study lasted about 60 minutes, and participants received a $10 honorarium.

Table 3.2: Experience of participants per interface.

Interface	**<1 year**	**1-3 years**	**4-6 years**	**7-9 years**	**9+ years**
Word	0	0	0	2	3
Facebook	0	1	0	3	1
Reader	0	0	2	1	2
Photoshop	1	2	2	0	0

The visuals of the interfaces used in the study were snapshots captured from the versions available in August 2018, using a Windows 10 PC with a 21.5-inch monitor. Participants drew sketches of the interfaces using pencil and paper. The study was approved by our local Ethics Review Board.

3.3.3.5 Procedure and Data Analysis

Before the study, we carried out an initial inspection of the four GUIs, where we saw that the layouts of Word, Facebook, and Photoshop differed from each other – often having multiple tabs or pages. Therefore, we decided to examine commonly used tabs and commands in our interviews. For example, there were 11 tabs in the Word interface, each with a slightly different layout. Therefore, we limited our study in Word to the *Home* and *Insert* tabs only. Similarly, from

Facebook, we chose the *Home* and *Profile* pages, and in Photoshop, we used Options bars (displaying additional tools for a tool selected in the left-side toolbar) for the *Move* and *Text* tools.

We ran a pilot study with four volunteers before the actual study to refine the questionnaires and tasks. During the study (carried out in a lab), participants first filled out a demographic questionnaire then proceeded to the interview. Each participant completed the study for only one interface, and all the sessions were audio-recorded. Later, we collected images of the actual interfaces used by participants in their workspace to cross-check their descriptions (this step was optional for the participants). We set out to answer three main questions:

1. Do users develop images of an interface: that is, do they develop spatial memory of commands and interface features, and remember their locations?

2. How strong and accurate are the images of an interface: that is, how accurately can users remember the locations of commands?

3. How do people use landmarks to remember commands and navigate in an interface, and what are those landmarks?

Our analysis began by transcribing the recorded interview sessions. We carried out a reflexive thematic analysis [36,37] in our study, particularly for the verbal descriptions and walkthrough tasks. Our study had several a priori goals, but we also wanted to be open to new ideas, so we coded the transcribed data using both deductive and open coding methods. While coding, we had some obvious a priori categories in mind – spatial references, landmark references, relative/absolute positioning, use of the overall frame of reference, difficulty in remembering a location, or ease of recalling a location. Then we generated potential themes for each of the four interfaces. For example, in Reader, we had codes, e.g., ‘abstract reference to the top’ and ‘vaguely pointing towards left;’ this resulted in ‘abstract direction’ as a potential theme. Additionally, we had ‘menu bar,’ ‘toolbar’ and ‘side panel’ as initial themes. Next, the first author used axial coding [57] to further refine the themes by checking their relations with the codes and discussed results regularly with other authors. The aforementioned potential themes, for instance, were merged into one theme, ‘interface layout’ for Reader. We repeated this hybrid deductive-inductive thematic analysis for the four interfaces to develop spatial images of the interfaces.

We also analyzed the sketches to check the sequences in which participants drew different elements and features of the interfaces. Furthermore, we analyzed and validated the accuracy of the verbal descriptions and pointing tasks by comparing them with actual interfaces (reported in Table 3.3).

3.3.4 Findings

The following sections present an overview of the mental images of the four interfaces, review their clarity, and report on the landmarks that contributed to building the images.

3.3.4.1 Cognitive Images of the Four Interfaces

We assessed the accuracy and coverage of participants' memory of the interfaces, using their verbal descriptions, their sketches of the interface, and their task walkthroughs. We found that people do develop images of the interfaces in their memory. The images we found are from a small population (although still reasonable [89]), and involve subjective responses; nevertheless, the data gathered from the study were rich and consistent enough to elicit meaningful, stable images of the four GUIs.

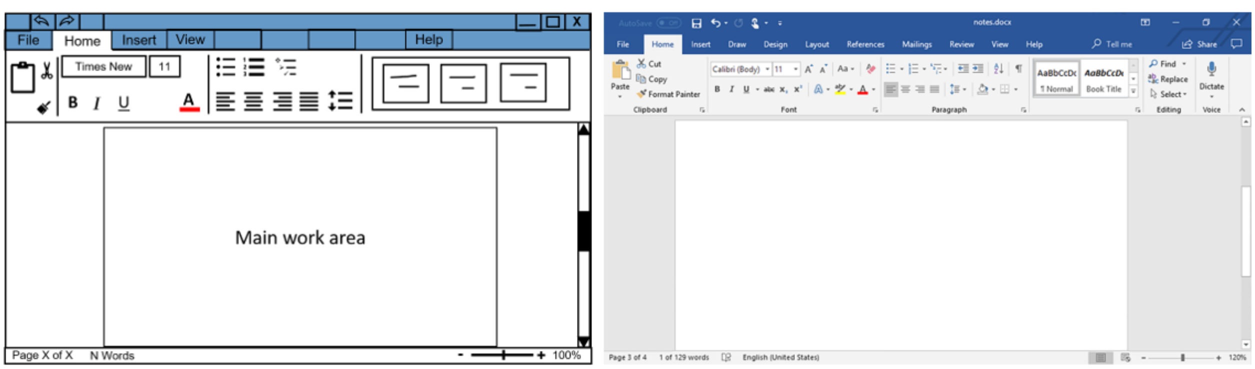

Figure 3.2: Images of MS Word's Home interface. Left: image reconstructed combining all participants' descriptions and sketches, right: a snapshot of the actual interface.

3.3.4.1.1 <u>Image of the Microsoft Word Interface</u>

The mental image of the Word interface contained both distinctive colour and structure: the *"blue colour menu"* at the top with a horizontal ribbon underneath and a *"white [coloured] page in the* *middle"* were two features everyone instantly remembered. To users, Word had *"lots of options [commands],"* yet was *"clean"* and *"organized"* because of its well-structured layout.

As shown in Figures 3.2 and 3.3, despite ambiguity in the name and order of tabs, everyone cited and drew a tabbed ribbon bar at the upper left corner. Although all interviewees knew the leftmost tabs (i.e., File, Home) and the rightmost (i.e., Help), surprisingly, none could recall the tabs from the middle (perhaps due to the serial-position effect [158]; further clarified in Discussion).

The most consistent element in the images was the ribbon, showing commands of a tab in groups isolated by vertical lines. However, we found contrasting images of the two ribbons: *Home* and *Insert.* The *Home* ribbon was the most striking of the two images. Everyone understood the two separate areas of this ribbon: the left half crowded with three groups of commands, and a large rectangular area, 'Styles' covering almost all the right half (see Figure 3.2). As one Word user described:

> All the tabs are at the top. In Home [ribbon], there are type of fonts […]. Then to the right of it, there are bullet points options. […] After that, there is a large box, styles.

Users also remembered the edges and boundaries of these groups, including their uniquely shaped icons (e.g., 4 out of 5 people described *"a box for the Font Name at the left").* They correctly recalled at least 20 of the 38 items from *Home*. Interestingly, we observed that participants first visualized the meaning of a command in their minds as they referred to the visual appearances (e.g., shape, or colour) of the corresponding icon while describing any command's location. For example, *"a painting brush"* was mentioned for the 'Format painter' command.

The *Insert* ribbon (see Figure 3.3), in contrast, had a slightly unclear image to the users, primarily because of the similarity in appearance among its commands and layout. This tab arranges nearly 32 icons, each with an upright rectangle icon, into ten groups, making it hard for users to identify the edges and boundaries of each group. As a result, users could name only 11 commands in this ribbon and were uncertain about their locations. However, 4/5 users correctly remembered the *"Equation and Symbol [commands] at far-right."*

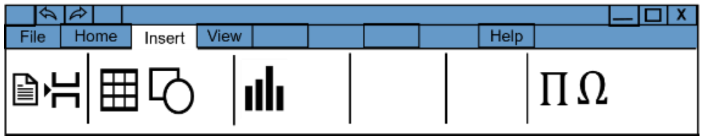

Figure 3.3: Image of MS Word's Insert ribbon interface reconstructed from all interviews (description and sketch).

Besides the ribbon, several items also consistently appeared in the cognitive images. For example, 3/5 people drew the 'Quick Access Toolbar' at the top-left corner; and everyone cited the correct locations of Pages, Word Count, and Zoom Level at the bottom (see Figure 3.2).

Overall, it appeared that the easy-to-recognize boundaries of command groups, uniqueness of the icon appearance, and icons' meaning helped users to construct the cognitive image of the Word interface.

3.3.4.1.2 <u>Image of the Facebook Interface</u>

A large number of commands in the Facebook UI and their distributed placement (often repeated – for example, the Messenger icon appears at three locations) contribute to a less clear image; participants visualized this UI it in several different ways (e.g., a *"blue and white"* GUI, *"lots of images," "News Feeds,"* or *"notifications")*. However, participants also recognized the structured layout with clear sections and meaningful commands that allowed them to form vivid spatial images of the *Home* and *Profile* pages (see Figures 3.4 and 3.5).

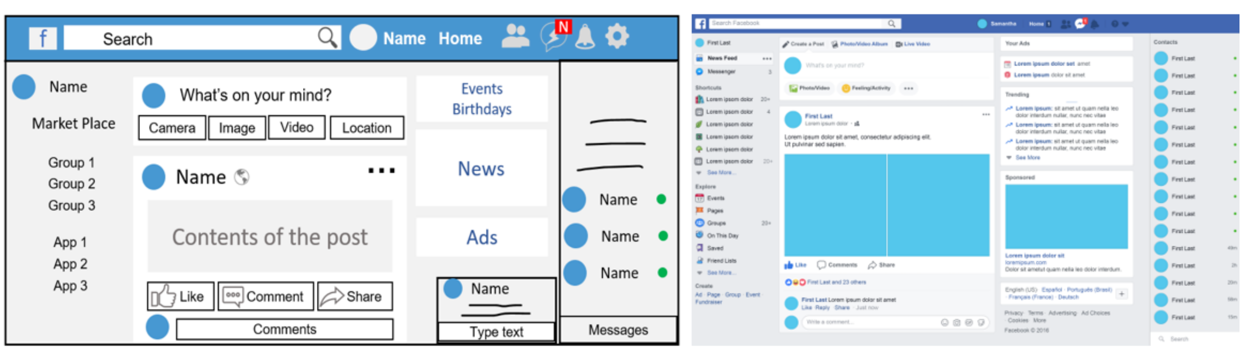

Figure 3.4: Images of Facebook's Home page. Left: image reconstructed combining all participants' descriptions and sketches, right: a snapshot of the actual interface.

We found that the three separate areas of the *Home* page appeared in all sketches and descriptions with clear boundaries (see Figure 3.4). The News Feed (NF), the widest among the three, sat in the middle, showing 'posts' in a scrollable stack, each with *"a unique box [a rectangular card showing all the elements of a post]."* All users identified the very first box as an *"area to create a new post with texts, photos, and videos."* Two users even quoted the text, *"What's on your mind?"* that usually appears inside that box. The rest of the boxes, displaying posts, had around 12 items in the layout, and everyone recalled (i.e., described and drew) at least 8 items at the correct locations. All stated that the left side's area held links to Groups and Apps; however, 3 out of the 5 users only vaguely recalled the contents as they seemed similar and lacked clear boundaries. The area at right, in contrast, held the *"Events, News, and Ads"* in vertically isolated areas that all participants identified correctly. Also, everyone mentioned the *"chat area at the bottom-right"* with correct details. One participant reported:

> Chat option is in the right-hand side of the screen. Normally it is hidden [collapsible menu] at the bottom, but clicking on it will show a list of people with green dots [availability] beside them.

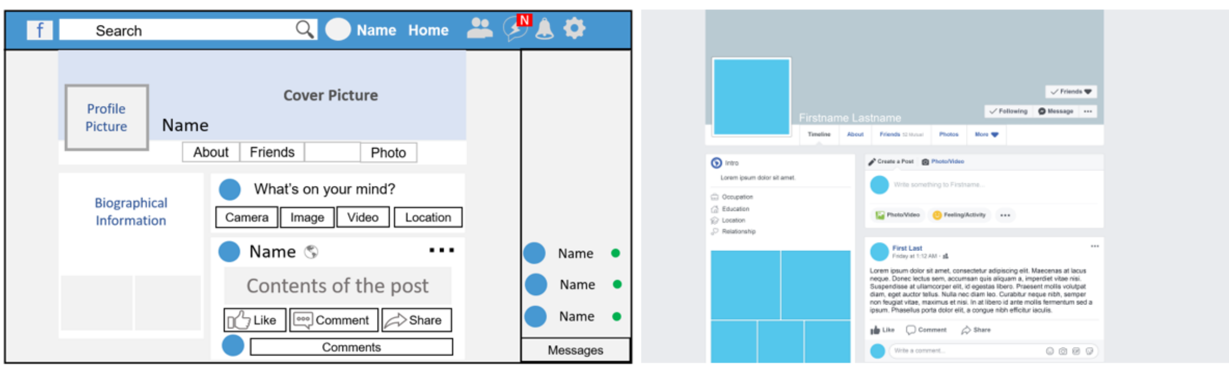

Figure 3.5: Images of Facebook's Profile page. Left: image reconstructed combining all participants' descriptions and sketches, right: a snapshot of the actual interface.

The mental images of the *Profile* page, on the contrary, had a large Cover Photo (CP) at the top, including a square Profile Photo (PP) at its bottom left (interestingly, one drew it as a circle), and a tabbed menu underneath the CP (see Figure 3.5). Though all reported these elements correctly, surprisingly, none recalled the five tabs correctly. The rest of the space was divided into two vertical areas, and almost all (4/5) noted that the large area at the right showed *"posts in boxes,"* similar to NF. We found that only two out of five users correctly recalled the sections for biographical information, photos, and friends located at the left, possibly because of their similar appearance (all appear in equal-sized rectangles).

There was, however, a vivid area in the topmost part of the two pages: the blue menu bar. As expected, all five users recognized a horizontal bar with two unique areas: a search box at the left and a few items at the right. The mental images were accurate: four users recalled at least 8 of the 10 items with the proper shape and order. Again, we observed a tendency among the users to describe the visual appearances of the command-icons along with their locations, similar to Word. One stated, *"There is a bubble [icon for Messenger] at the left to the bell [icon for Notifications]."* Another user said, *"a thumbs-up at the bottom left corner,"* when describing the 'Like' command.

3.3.4.1.3 Image of the Abode Acrobat Reader Interface

The task of 'reading a PDF' is so prominent in Reader that everyone remembered it first, even before recalling other items: the *"red logo," "colourful,"* yet *"annoying sidebar."* We found consistency in the mental images that included a PDF document pane in the middle and commands around its three sides, a toolbar above the PDF document, and two side panels. Surprisingly, despite being the simplest GUI among the four, participants struggled more to recall command locations in Reader, possibly because of the linear placement of commands in the toolbar (see Figure 3.6) without clear grouping.

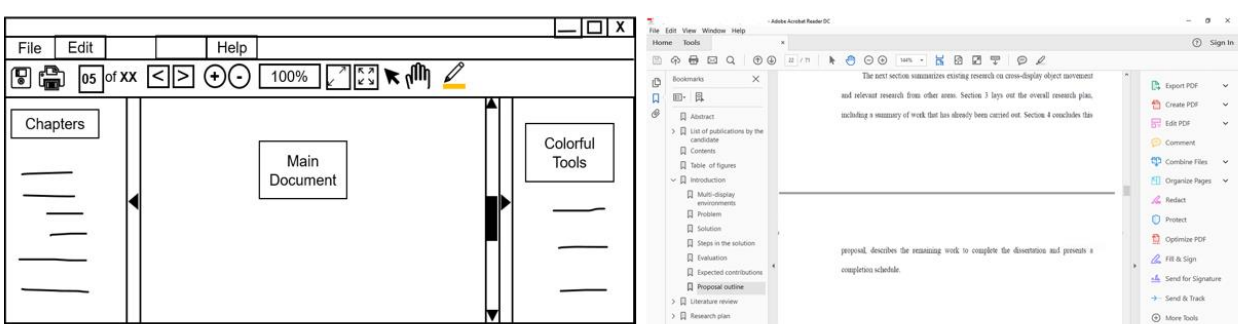

Figure 3.6: Images of Adobe Acrobat Reader interface. Left: image reconstructed combining all participants' descriptions and sketches, right: a snapshot of the actual interface.

As seen in Figure 3.6, all users mentioned the tabbed menu at the top, but similar to Word and Facebook, none could exactly recall exact names and order. They were, however, unanimous that a horizontal toolbar was present just above the PDF. Interestingly, though nearly all (4/5) users could name at least 12 of the 19 commands linearly placed in the toolbar, they could not precisely remember their locations. Also, despite uncertainty in the direction (i.e., vertical or horizontal), all knew the *"two arrows [icons for Next page and Previous page]"* to navigate pages, and three people were confident that a *"small square for page number"* was beside those arrows. However, four of five users recalled the Highlight tool precisely at *"far down the right"* and described the visual appearance of its icon, portraying a *"pen slightly tilted with colour below it, usually yellow."*

Apart from the toolbar, two side panels always appeared in all the images, but most users (3/5) kept those hidden to make room for the PDF. Though users could imagine the *"Table of Contents"* in the panel at the left, none recalled the *"colourful"* tools accurately from the panel at the right.

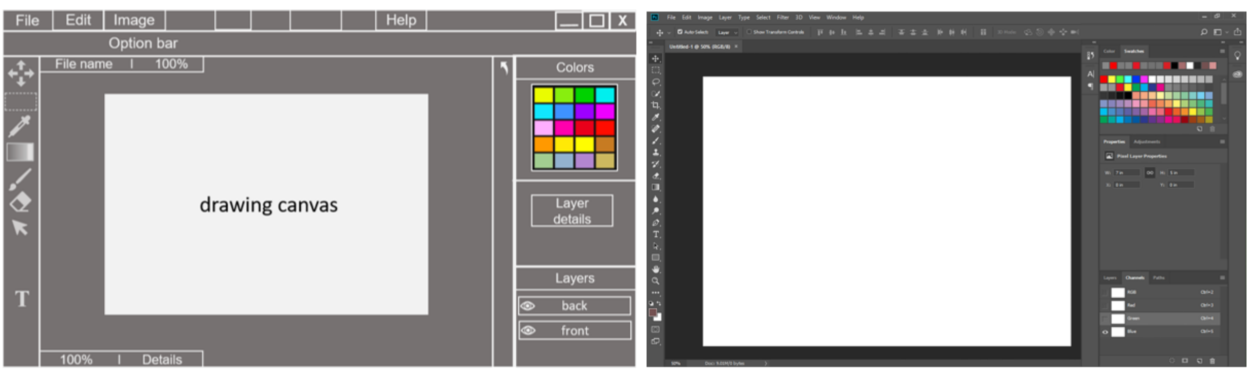

Figure 3.7: Images of Adobe Photoshop interface. Left: reconstructed combining all participants' descriptions and sketches, right: a snapshot of the actual interface.

3.3.4.1.4 Image of the Adobe Photoshop Interface

Participants visualized Photoshop as a *"black-themed"* GUI consisting of *"lots of tools."* Similar to Reader and Word, a large canvas at the middle occupied a significant part of all images, providing a strong anchor for three nearby areas: the Tool Panel at left, the Options Bar at the top, and a Side Panel at right (see Figure 3.7). Overall, it was evident that everyone remembered the overall layout of Photoshop; however, the linear presentation of commands without a clear sense of division made the interface less memorable.

The Tool Panel was strongly visible in all the cognitive images; however, we found ambiguity in the name and order of commands. Though all recognized this vertical panel, participants could only name about 10 of the 23 tools, and only vaguely recalled their locations. However, similar to the previous three GUIs, users referred to the visual appearances of commands while attempting to recall their locations. As one said regarding the 'eraser' tool:

> It [the eraser command] is at the left side panel, kind of in the middle and looks like a physical eraser, [placed] in angle [slightly tilted] with a shadow [underneath].

We found the Options Bar located above the canvas even less memorable, as its tools (and layout) change based on the item selected in the left panel. Only two people mentioned that *"it shows additional tools,"* and surprisingly, none correctly recalled any tool from that panel. The right-hand panel, however, was more vivid because of the three clear sections. Almost all (4/5) users correctly recalled the two sections: Colors at the top and Layers at the bottom. One user correctly remembered the History tool beside the right panel.

In summary, we found that users developed images of all the interfaces in their minds, but the memorability of the interfaces varied: people could recall the locations of commands more reliably in Word and Facebook than Reader and Photoshop (reported in Table 3.3). Despite the differences, it was evident that users developed spatial memory of commands in all the four GUIs, especially in areas having clear frames of reference such as corners and edges in the GUI environment, and command groups with easily identifiable boundaries – i.e., areas with landmarks.

3.3.4.2 Accuracy of the Cognitive Images

We analyzed the location data collected from the study to assess the images' accuracy (Table 3.3). Participants were asked to verbally describe locations and nearby commands of five commands they frequently used. For the verbal description and nearby commands, when users accurately described an item's location in a GUI and at least one adjacent item respectively, we treated it as *correct* (e.g., in Word, *"Symbol is at the right edge of the list, in the Insert ribbon"*), otherwise *incorrect*. However, when they misdescribed a location but were close or used a general direction (e.g., *"at the top"*), we noted it as *partially correct*. For the sketch and washed-out images, *correct* indicates that the pointing gesture was within 10mm, with larger errors marked as *incorrect*.

We report the number of events (e.g., correct, partly correct, or incorrect) for the four tasks across all the participants (percentages denoted in brackets) in Table 3.3. Overall, locations for both Facebook and Word were more accurate (at least 76%) than Reader and Photoshop (less than 64%). In participants' verbal descriptions, 90 of the 100 locations were nearly correct or correct. For Word and Facebook, people were more accurate (over 76%) than for Reader and Photoshop (below 44%). For nearby commands, among 100 cases, over 57 were correctly recalled. Again, Reader and Photoshop had lower accuracy (less than 36% correct responses).

Table 3.3: Results (in frequencies; percentages denoted in brackets) of the four tasks used in the study.

	Verbal Descriptions			**Nearby Commands**			**Sketch**		**Washed-out**	
GUIs	Correct	Partly Correct	In-correct	Correct	In-correct	No Reply	Correct	In-correct	Correct	In-correct
Word	19(76%)	5(20%)	1(4%)	19(76%)	6(24%)	0(0)	21(84%)	4(16%)	21(84%)	4(16%)
Facebook	23(92%)	1(4%)	1(4%)	21(84%)	4(16%)	0(0)	24(96%)	1(4%)	24(96%)	1(4%)
Reader	10(40%)	12(48%)	3(12%)	8(32%)	14(56%)	3(12%)	13(52%)	12(48%)	13(52%)	12(48%)
Photoshop	11(44%)	9(36%)	5(20%)	9(36%)	13(52%)	3(12%)	15(60%)	10(40%)	16(64%)	9(36%)
Total	63(63%)	27(27%)	10(10%)	57(57%)	37(37%)	6(6%)	73(73%)	27(27%)	74(74%)	26(26%)

However, total accuracy improved in both the sketch and the washed-out interfaces (over 73% accuracy), because the available layout outlines, group borders, and shapes provided better landmarks to recall commands. The accuracy of pointing to command locations in Word and Facebook was at or above 84% correct, while Reader and Photoshop had below 64% accuracy, suggesting their weaker support in developing spatial memory of commands (discussed further below). Overall, these results are an indication that we successfully elicited the GUI images from users' minds.

3.3.4.3 Landmarks in the Graphical Interfaces

Although the individual spatial images of the four interfaces seemed sparse and varied in layout and number of commands, clear patterns emerged when we analyzed them together. Based on Lynch's notion of landmarks [143], we looked for unique, stable, always-visible features and structural elements (e.g., layout) in GUIs, including elements from both inside (e.g., groups) and outside (e.g., device corners) the interfaces that participants used as landmarks to remember the locations of commands in the four GUIs. There are, of course, other factors (e.g., subjective cognitive ability) that can influence the formation of interfaces' spatial images; here, we mainly focus on the role that landmarks play.

We identified four different landmark types from our data (see Table 3.4). Although these landmarks acted together to form the spatial images, for simplicity, we present them separately. The following sections present these landmarks and report how users employed them to recall commands.

Table 3.4: Types of landmarks present in the interfaces and their functions.

Type of landmark	Example	Function
Interface Layout	Structure of a GUI containing commands and other elements: toolbar, ribbon, side panels, etc.	Provides an abstract location information of commands and other elements.
Command Group	Font and paragraph groups in Word's Home tab. A card displaying a post in Facebook's News Feed.	Provides an absolute location information of a command relative to other commands of the group.
Corner and Edge	Corners and edges of a GUI window, or a group (Facebook's card).	Provide a clear frame of reference for the location of a command.

Type of landmark	Example	Function
Icon Visual Appearance	Format painter icon looks like a “painting brush”.	Create memory anchors for a command by encoding its location information to its visual appearance.

3.3.4.3.1 Interface Layout

The layout is the overall structure of an interface that arranges commands in menus, toolbars, and panels, placing them around the main work area. Users traverse through different parts of a GUI’s layout in order to interact with its commands. It allows users to form a birds-eye-view image of the interface, helping users remember the locations of commands and navigate to them. Therefore, an interface’s layout can act as a “reference frame” type of landmark. However, layouts do not reveal too many details about the locations of commands and other elements, because each region of a layout usually contains several commands. We found that all participants understood the unique layouts of the four GUIs and used them as a general reference frame for remembering relative locations of commands. As a result, even participants who were uncertain about the location of some commands had a clear idea of how the elements were laid out in the interfaces. For example, one Word user said:

> All the tools [commands] are located at the top and grouped into several tabs such as File, Home based on item category [the ribbon toolbar]. Items related to texts [editing] are in the left [Home tab].

From our analyses of the sketches and interviews, it was evident that the layout of an interface came to users’ minds first when they tried to visualize any command. Similar to the *paths* described by Lynch [143,222], the layout connected different regions of an interface, providing routes for users to navigate through the GUI easily. During the sketch task, 17 out of the 20 users drew the layout first and then they placed commands in different parts of the layout, allowing users to visualize the complete picture of an interface. A Facebook user described the Home page:

> It has three separate areas with the ratio 20:60:20. The left area shows the ‘Groups’ […], middle one displays the posts, images, and videos; and the right side is for events and news updates.

Computer applications are primarily designed for specific tasks (e.g., *Word* for writing documents, *Photoshop* for editing images). Users observe and interact with the commands of an interface while performing the task at a specific area in its layout – the main workspace. We found this large area of an interface's layout (e.g., "a blank white page" for Word) serving as the primary point of focus in our four GUIs, forming a vital landmark [200] that all participants referred to remember commands and others areas of the interfaces. One user recalled various parts of the Photoshop interface using the main workspace as a landmark:

> There is a big canvas in the middle, where we can draw or edit. On its left, there are some useful tools for the work. It has some tools in a menu at the top as well. In the right, there is a panel of colours and layers.

We also found that in the walkthrough task, 15 of the 20 users relied heavily on this central area of an interface's layout to describe locations of commands and other areas, and interestingly, it was the starting point of their narration. For example, when we asked people to describe step by step how to search for something on Facebook, a user replied:

> You see the large area in the middle [News Feed] with lots of posts and pictures? Right above it, there is a blue bar with a white box in its middle. You can type there whatever you want to find.

Our analyses further revealed that people used general spatial directions – a widely used real-life reference frame – inside an interface's layout to provide coarse location knowledge of commands. For example, one Word user said, *"The Bulleted List command is located at the top [of the layout]."* This answer, although under-specified, is not incorrect – it pointed to the vicinity of the actual location and eventually helped the user to find the command. The majority (14/20) of the users (across all interfaces) used terms like *"towards the left"* or *"to the bottom"* area of the layouts in their responses to recall the locations of a few commands vaguely (partly-correct responses in Table 3.3), mostly for interfaces (e.g., Reader and Photoshop) and particular areas of Word (e.g., the Insert ribbon) that lacked adequate landmarks. One Photoshop user even mentally divided the Tool Panel into four quarters and used them as references: *"a quarter down from the top"* and *"3 quarters down."*

3.3.4.3.2 Command Group

Graphical interfaces usually arrange commands into several tabs or sections based on utility or similarity [211,241], and even within an individual tab/section, sub-groupings of commands with specific areas are evident. These semantically divided groups can help users find commands, particularly when they are novices. However, expert users can recall the locations of commands using the *visual* representations of the grouping as landmarks (rather than the semantics). We found that the users in our study relied on the visual command groups when visualizing an interface. One Word user described the Home ribbon using the grouping feature as a landmark:

> In the Home [ribbon], there are types of fonts, change size and B, I, U [icons for Bold, Italic and Underline]. Then right to it [that group], there are bullet-point and alignment options. After that, there is a large box for styles [another group].

Usually, UI groups are denoted with clear borders (e.g., vertical lines in Word and boxes in Facebook) that most users (15/20) drew in the sketches. Similar to Lynch's concept of a *district* [143,222], users could even cognitively go inside a group having distinct features and layouts, and visualize its commands accurately. For example, each post on Facebook appeared as a group in a rectangular box that everyone recalled clearly. As one user said:

> It [each card/box in the News Feed] shows an image in a circle and [post creator's] name at the top left. […] Below, there are buttons for Like, Comment and Share. After that, it has other people's comments.

Our analyses showed that in at least 76% of cases, users correctly recalled nearby commands for both Word and Facebook, compared to less than 36% for both Reader and Photoshop (see Table 3.3). These differences are not due to different experiences with the UIs, as all of our users were highly experienced (see Table 3.2). One reason specific to Reader could be that people do not use commands as often in Reader as they do in Word – people mostly just scroll and read, so they might have less familiarity with the commands (this is certainly not the case for Photoshop). A further reason that is related to the design of the GUIs is that the linear placement of commands with no group-based arrangement in Reader and Photoshop reduced memorability. Our results indicate that the clear and easily separable command grouping served as a reliable landmark in painting vivid spatial images of the interfaces, and the absence of landmarks made GUIs (at least parts of them) less memorable.

Interestingly, we further noticed that when users identified a command from a group, they could quickly recall its nearby commands. Also, 17 of the 20 people used the relative association among the commands to devise informal groups and used the location of a distinctive command as an anchor to recall the locations of adjacent commands. One Facebook user said, *"[...] a bubble [icon for Messenger] at the left to the bell [icon for notification]."* Similarly, to describe the location of the page navigation arrows in Reader, one user relied on relative position:

> In the menu, there is a small box with a number on it [page number]. The two arrows [for navigating pages: previous and next] are right next to it.

In both cases, the references were accurate, but to find the actual locations, users must locate those anchor commands first, or else the landmarks become invalid.

3.3.4.3.3 Corners and Edges

The corners and edges available both outside and inside digital interfaces are an obvious set of landmarks [95,218] that we found people explicitly brought up in the study.

Corners

GUIs usually appear inside physical screens having four clear corners. These external corners are always visible to computer users and can substantially aid spatial learning and navigation for the commands appearing near those locations [208,218] (though only useful when a GUI is maximized on the screen), similar to Lynch's [143] *point landmarks* in cities. Also, the corners of a GUI window can provide a strong landmarking facility. In our study, we saw users strongly relying on these landmarks to recall commands correctly in all four interfaces. For example, everyone drew the 'Close,' 'Maximize,' and 'Minimize' commands at the top-right corner of all the sketches. All of the Word users recalled that *"zoom level is at the bottom-right corner,"* and the *"number of pages is at the bottom-left [corner]."* All the Facebook users mentioned that *"Messenger is at the bottom right corner."* Even in Photoshop, where the mental images were less accurate, almost all (4/5) users correctly recalled the 'Magnification status' in the bottom-left corner of the window (see Figure 3.7).

Corners, on the other hand, are also available inside GUIs. These internal corners, however, are not as prominent as external corners, but exist within the layout of an interface as well as in clearly marked command groups. Similar to external corners, internal corners can provide stable anchors for commands located near them. In the study, 13 of 20 users explicitly referred to these landmarks while remembering commands. For example, a Facebook user said, *"The Like [button] is at the bottom-left part of the box [displaying the post (see Figures 3.4 and 3.5)]."* Another user made use of internal corners in the walkthrough task while describing the way to delete a post:

> Look at the middle [in News Feed]; a box is showing the post [you want to delete]. Go to the top right corner of the box; there are three dots [icon for options].

Edges

Similar to corners, external screen bezels and edges have spatial features (i.e., they are stable and easily visible) that help users in developing memory of commands [191,218]. Following Lynch's [143] definition of edges in physical spaces, we found that the four external bezels (when a window was maximized), as well as the edges of the GUI window, provided spatial anchors for the commands located near them, at least at an abstract level. However, in our study, users did not differentiate between the bezels of the actual screen and the window of a GUI. References to these landmarks, e.g., *"near the right-side"* or *"at the top-side,"* repeatedly appeared in our analyses, especially for the Photoshop and Reader interfaces. These landmarks, however, did not refer to a specific command or an absolute location; instead, they pointed to a *command group* of a particular GUI region. In other words, these edges served as "signposts," providing abstract directional information for other elements in GUIs. For example, one Reader user said: *"There are some commands at the right [near the bezel]."* Another Photoshop user relied on this landmark to correctly visualize the Tool Panel:

> [The] most important items are at the left side [near the left bezel] of the screen.

Internal edges, in contrast, provided explicit references to the specific locations of commands. The internal edges refer to the boundary of a command group, or the starting and ending points of a linear command group. Strong dependence on the internal edges was apparent among users in our analyses. For example, as seen in Figure 3.8, one Word user said,

> The 'Insert Symbol' is at the right end of the Insert Tab. Not to the far right, but at the end of that list.

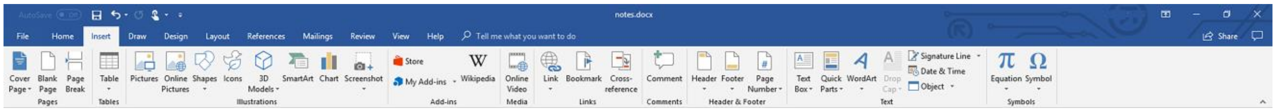

Figure 3.8: A snapshot of MS Word's Insert ribbon.

Although participants used internal edges less frequently than corners, they made use of these landmarks to visualize at least part of the horizontal tabbed menu bar correctly in all GUIs. We also found that participants faced more difficulty recalling locations, particularly in the middle of the menu bars where edges were weakly marked. For instance, while sketching the Reader interface, most of the users (3/5) drew the menu bar with 4-5 blank tabs, but only wrote 'File' in the first tab and 'Help' in the last (see Figure 3.6).

3.3.4.3.4 Icon Visual Appearance

Graphical interfaces use two-dimensional icons to portray the underlying meaning of the corresponding commands visually. Shape, size, and colour are three visual attributes that make icons meaningful [160], and usually, novice users rely on these features to find an icon (i.e., a command) in a GUI [50,187]. Although these visual features do not provide any information about a command's location, interestingly, we found that participants frequently referred to them while recalling any command's location. In most cases, participants imagined the visual appearance of an icon first before describing its location – indicating, people might have encoded icons' spatial information into its visual features. For example, one Word user recalled 'Insert table' as *"a square icon with a couple of rows and columns located in the left part of the Insert [ribbon]."* Another cited 'Line spacing' command using the visual features and meaning of its icon as a referencing mechanism – landmark:

> […] two up and down arrows with two/three lines [icon for Line spacing], near the centre of the Home ribbon, beside the Alignment [commands].

We also observed that all participants could clearly visualize the shape of an icon while recalling it. The shapes were clear enough in users' minds that they drew those shapes accurately in the sketches, even though they were often incorrect about their locations. For example, although people rarely recalled tools from the Insert ribbon in Word, nearly all users (4/5) identified the exact location of the 'Insert Symbol' as a *"round Greek letter at the end."* Also, a user drew a hand to represent the 'Pan' tool of Reader (see Figure 3.6). All Facebook users drew a thumbs-up for the 'Like' button and a circle at the top-left corner of a post that usually shows the image of the post creator (see Figure 3.4). Often these shapes resemble real-life objects that people recognized and exploited to remember commands. For example, one Photoshop user mentioned the *"dropper we use in the chemistry lab"* while recalling the 'Colour Picker.'

In addition to shape, we found that the colour of a command was another prominent feature that users often used as a landmark. Colour has a unique ability to make one icon easily separable from others [147,200], and we saw that people recognized it and used it as a reference point to recall those items, along with other features. All the Facebook users, for instance, mentioned:

> [There is] a blue bar at the top [of the interface]. [It has] a long and white colour rectangular box [search bar] with a magnifying tool.

Three of them also mentioned changes in the icons – e.g., one user remembered *"a number in a little red box"* when a notification is received (see Figure 3.4). Sometimes two visual features act together to make the images more vivid. One participant described the 'Font Color' command in Word as *"a big 'A' and a red colour line under it."*

Our analyses indicated that the size of commands in an interface also provided clues for spatial memory. We saw that users recognized the size differences in some commands, and later that helped them to identify commands correctly. For example, in the Home ribbon of Word, there were two small but differently-sized boxes in the Font group, and almost all users (4/5) correctly identified the relatively large one as the 'Font Name' command because of its size.

Overall, we found that the frequent users of these four GUIs developed spatial memory of commands by exploiting the available landmarks in the interfaces, and that they heavily relied on those landmarks to retrieve the command locations from memory.

3.3.4.4 Ambiguities in the Cognitive Images

The individual spatial images of an interface overlapped and collectively formed a 'composite image' [143] of the interface. However, in our analyses, we observed some uncertainties in the images, at least in some parts of them, at two levels: individual images and composite images.

3.3.4.4.1 Ambiguities in Individual Images

Individual-level ambiguities occurred in isolated images and did not influence the composite cognitive image of an interface. For example, we found confusion among Facebook users with the location of the ‘Messenger’ command – it appeared at two locations: at the bottom-right and top-right. Interestingly, both were correct. Messenger does appear at two locations, as mentioned by the users. So, when we asked about its location, three users reported only one (either top or bottom), and two mentioned both locations. In another instance, although the ‘Profile Photo’ command in Facebook is squared-shaped and appeared at the bottom-left part of the Cover Photo, one user drew it as a circle and placed it in the middle. The same user mentioned a ‘Market place’ option in the top-left part of the News Feed, whereas it was absent in other (4/5) images. A probable reason could be that when companies redesign their UIs, it may confuse users’ mental images and break their landmarks (elaborated in Discussion).

Similar events occurred in Word images as well. For example, although most users mentioned the ‘Font Color’ icon (a big ‘A’) was in the bottom-right corner of the ‘Fonts’ group in the Home ribbon, interestingly, one user recalled it at the middle bottom-centre area (perhaps different screen size was responsible; discussed further below).

3.3.4.4.2 Ambiguities in Composite Images

We noticed some uncertainties were consistent in most of the mental images, irrespective of individual differences. There were certain areas in the interface images that were less memorable among the users. For example, the tabbed menu bar in Word, Reader, Photoshop, and Facebook (Profile Page) interfaces was one place where all users struggled. Although everyone recognized the menu with linearly placed commands, surprisingly, none could recall the names of tabs and their order accurately (could be due to the serial-position effect [66,158]), except for the first and last tabs (see Figures 3.2, 3.3, 3.6, and 3.7). As one Word user said:

> I know there are several tabs after Home and Insert that I often use, but seriously,
> I cannot remember their names! I have even used it this morning.

Apart from the tabbed menus, there were certain general areas in the composite cognitive images of individual interfaces where most users faced difficulty in recalling commands accurately: the

Insert ribbon in Word, Reader's toolbar, and Photoshop's tool panel (left-side). Although everyone in our study knew where those menus and toolbars were located and could name some of their commands, they could not recall the locations of those commands accurately. As a result, one Photoshop user described the tool panel and toolbar vaguely:

> On the left-hand side of the image [main workspace], there are some tools [tool panel] for work. [And] like other apps, there is a menu [toolbar] at the top part of the screen.

Overall, despite these minor ambiguities, the mental images we elicited were vivid and accurate enough to identify available landmarks in them. We detected four landmarks that supported users to develop spatial memory of commands and found that people heavily relied on those landmarks to recall the commands later.

3.3.5 Discussion

The key findings of our study are:

1. Frequent users of graphical interfaces do develop vivid spatial images of those GUIs in their minds, and we were able to successfully evoke those images;

2. Users rely heavily on landmarks to familiarize themselves with GUIs and remember the locations of commands;

3. We identified four different landmarks that aid in developing spatial memory of the interfaces.

The following sections reveal more insights into our findings and present future directions, along with design implications.

3.3.5.1 Interpreting the Results

3.3.5.1.1 <u>Reliability of the Interfaces' Images</u>

The images of the interfaces we generated from the study were rich in information, revealing how people perceived the visual forms of those GUIs and leveraged the landmarks readily available in a GUI environment to form spatial memory of commands. Although we interviewed only 20 people (5 people for each of the four GUIs), we believe the following two reasons indicate that our evoked interface images were of good quality. First, the images of the GUIs people had in mind were the collection of experiences they had with those GUIs [98]. Since our participants were reasonably experienced and frequent users of those four GUIs (see Table 3.2), they could compose vivid and detailed images of the GUIs [195]. In our analyses, we found several statements that only an experienced person could share. For example, one Reader user described the 'Search' command as *"A blank white box with 'Find' and a magnifying glass pops up at the top-right part."* Another Word user recalled, *"Temporary tabs will appear if you select a table or an image!"*

Second, we compared these evoked images to their real counterparts to check their accuracy (see Table 3.3). Also, we reviewed snapshots of the GUIs from participants' workstations (an optional part of the study) that participants sent after the study. Although the images did not cover everything that appeared in the actual interfaces, we found that they included the basic structures and visual features along with the frequently used commands. Overall, we managed to elicit clear images of the four interfaces, stable enough to be treated as the 'public images' [143] of the respective GUIs.

3.3.5.1.2 Icons' Visual Appearances Acted as Landmarks

Among the four landmarks we identified, although the visual appearance of a command did not disclose any spatial information, interestingly, everyone actively referred to it while recalling its location. One of the main reasons for this finding can be explained by the stages of learning (cognitive, associative, and autonomous) [14,77] and forming spatial memory [195]. Users in a new GUI, particularly novices (i.e., those at the *cognitive stage*), find an icon/command using visual features that represent its meaning, often attaching personal stories or experiences to its visuals [45,58] as part of the development of landmark knowledge [74]. After locating a command, they encode the location information along with its visual appearance (obtaining route knowledge [209]) and eventually progress to the *associative stage*. When users become experts (i.e., the *autonomous stage*), they can simply recall the location of commands from their memory, using their acquired survey knowledge [195]. Since the visual appearance of a command was involved at the very early stage of learning to recall its location, participants in our study might have imagined that appearance first – indicating visual appearance was a potential landmark.

Another reason could be the recognizability and memorability of icons. People can recognize and locate commands more effectively when the icons representing the commands portray accurate meaning [33,45]. Research also suggests that some icons or images are more easily memorable than others [108,178,202], perhaps because of personal experience [58] or the presence of landmarks – an avenue we will explore in future. In our analyses, we saw that users could make use of three landmark types (GUI layout, command group, and corner and edge) to reach the vicinity of a command's location; however, it could be the visual appearance of the command that they imagined in their minds (and used as an anchor) to reach its actual location.

3.3.5.1.3 Ambiguities in the Images of the GUIs

Some degree of differences in individual images of the GUIs were inevitable; in fact, these contrasts made the composite images more informative. However, we observed some unexpected distortions in both individual and composite images of the GUIs (reported in section 3.3.4.4). Four main reasons could account for these confusions – first, the difference in screens' sizes forces an interface to rearrange its elements. As shown in Figure 3.9, Word rearranges all the items under a group from a two-row layout (Figure 3.9: top) to a three-row layout (bottom) layout when the window is resized from large to small. As a result, two people coming from these two different instances would produce two slightly different images of the same GUI: 'Font Color' command (letter 'A' with a red line underneath) at the bottom-right corner in Font group (Figure 3.9: top) in one image, and the bottom-centre in another (bottom).

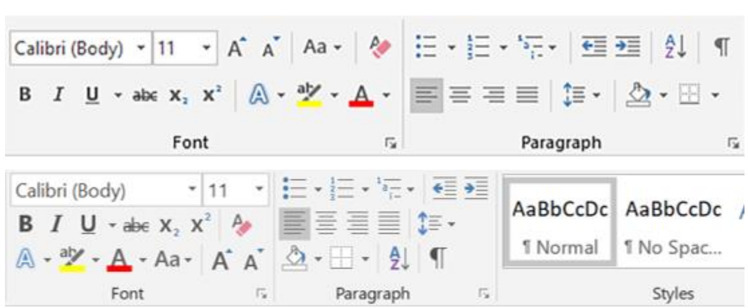

Figure 3.9: Snapshots of MS Word's Home ribbon from two differently-sized screens. Top: 22-inch desktop, Bottom: 13.1-inch laptop.

Second, interfaces often change over time because of updates that add or remove items and alter layouts [242], which, in most cases, do not reach all users at once [243]. These changes force users to adjust their images and could lead to some ambiguities in the individual images of a GUI. However, we noticed that users were aware of these temporal changes of the spatial images: as one Facebook user said, *"There is a 'Market Place' at the left side of the News Feed, under my name. It was not there before!"* Still, modification of the interfaces over time is inevitable, but care should be taken that it does not force users to shift the image of an interface that was developed over a long period.

Third, the lack of adequate landmarks in certain regions of the GUIs could have contributed to the uncertainties found in their composite images. For example, the tabbed toolbars in all the images were unclear, along with Word's Insert ribbon, Reader's toolbar, and Photoshop's tool panel. However, a common trend we noticed in those areas that the commands were linearly placed without arranging icons in clear groups or with easily distinguishable boundaries. As a result, people could not associate those commands to memory anchors, and later struggled while retrieving the locations from memory. A potential solution to this problem is using artificial landmarks [218,222] in those areas to improve memorability.

Last, our participants differed in several ways: ages ranged from 18 to 43; we had both men and women users; participants varied in terms of both the total time using the GUIs and the weekly frequency; and users' backgrounds also varied. These demographic differences may have led to some ambiguities in learning and remembering command locations in GUIs. For example, due to extensive usage, a professional Photoshop user described an aspect of the GUI ("Left side panel can be customized based on your need") that other users did not know about. Research also suggests that people's ability to recognize and successfully employ landmarks to recall locations and navigate an environment can vary depending on age and gender [74,136]. The small sample size in our study has prevented us from investigating landmark usage in GUIs based on participants' age and gender – an avenue we plan to explore in the future.

3.3.5.2 Generalizing the Results and Design Implications

Our study reveals that the readily available landmarks in a GUI help users orient themselves to its commands and can aid them in developing spatial memory of the interface. The benefit of landmarks in building spatial memory of commands and expertise with GUIs is not unknown – in fact, several HCI researchers have already demonstrated the potential of landmarks through prototypes with either external objects (e.g., hands) or strategically created items (e.g., images and colour blocks) [79,80,154,216,218,219]. Our work fundamentally differs from these previous studies in that we specifically focused on expertise development in standard GUIs present in the real life and which were not augmented with landmarks. To the best of our knowledge, this is the first work that explored the basic questions of what contributes to the development of expertise in standard GUIs, and what role landmarks play in this expertise.

Our findings help to explain why frequent users of a GUI, empowered by the developed vivid spatial images of the interface, can recall command locations easily from their memory; our results agree with existing research findings [50,92,183,218]. However, we also encountered several instances (about 27% in all tasks) where even our expert users struggled to recall command locations accurately – indicating some regions in the GUIs (e.g., tabbed menu bar; see section 3.3.4.4) did not provide enough cues for memory (i.e., landmarks) to build a clear spatial memory. Therefore, designers of GUIs should consider including more landmarks in an interface to support better recall, even for experts.

- **Implication 1:** Include landmarks in GUIs to support better recall.

This implication appears to contradict a widely practiced usability heuristic – 'recognition rather than recall' [161,244]. We acknowledge that recognition is a vital part of spatial memory development [170,208], and particularly novice users rely on recognition (i.e., visual search) to find the location of a command. However, when novices begin transitioning to experts, they tend to recall locations from memory. After becoming experts, the goal is a quick recall than a slow recognition [50]. Therefore, interfaces should be designed in a way that would facilitate both recognition and recall, so that they can support the user's transition from novice to expert [51,186].

- **Implication 2:** Interfaces should facilitate both recognition and recall.

In addition, results indicate four different types of landmarks that help users learn and recall the locations of commands in GUIs. Interestingly, the landmarks we identified (e.g., GUI layout, command group, and corners) are already present in the GUI environment, and designers include them in GUIs as a part of standard design (e.g., Gestalt principles) [211,241] and usability practices such as Nielsen's design heuristics [161,162,244], but not as landmarks. The novelty of our findings is that we have discovered an additional value of these already-existing design elements – i.e., their value as landmarks. Therefore, designers can consider using those GUI elements and features (e.g., the visuals of icons) as potential landmarks in order to design more easily memorable GUIs, and incorporate the idea of landmarks in current design practices.

- **Implication 3:** Incorporate the 'idea of landmarks' in existing design practices.

Another interesting finding is the use of icons' visuals as landmarks to remember commands. Although the value of visual appearance to differentiate and remember commands is known [33,45], no other work has provided evidence of using visuals as a promising spatial referencing mechanism (i.e., landmark) that can help users to develop spatial memory for commands and later enable memory-based recall. However, further research is needed to understand more about how the appearance of an icon can represent spatial information or how designers can exploit it to design more memorable GUIs. Besides these four landmarks, designers can consider other useful landmarks [200,218,222] (e.g., images, colour blocks), but care should be taken so that they do not become distracting [175].

- **Implication 4:** GUI layout, command group, corners and edges, and visuals of icons can be useful landmarks.

3.3.5.3 Limitations and Future Work

We see three main limitations in our study that prevent us from generalizing the findings more broadly. First, the sample size we used was small, though reasonable enough to generate stable trends [89] that created the basis for our analyses. Second, the cognitive images that users provided to us could be influenced by their subjective experiences (which is inevitable in qualitative studies [143] but should be followed up with a larger-scale study). Nevertheless, we believe the users' vivid images enabled us to see a reasonably complete picture of an interface. Last, we limited our investigation only to desktop interfaces involving command selections, while landmarks can be valuable in other spatial tasks (e.g., interaction with 1D documents: video and text [7,49,219], results sets in visual workspaces, or big-data visualization systems); this choice was made for simplicity and to avoid introducing an additional variable. We plan to investigate spatial learning in alternate structured visual information in future. Despite this limitation, the four interfaces we chose were diverse, and the data we gathered showed this diversity.

There are several ways we can continue our research in the future. First, people use GUIs in multiple platforms that often vary in size, which can force users to deploy multiple instances of the same GUI. We plan to carry out a large-scale study to generate public images of GUIs in users' native environments: desktops, smartphones, and VR and AR. Second, data-driven or adaptable interfaces often place frequently used commands in convenient locations, but in doing so, they may complicate the natural process of developing spatial memory [183] (particularly if a recency metric is used rather than frequency). We plan to investigate this issue in future (e.g., what frequent users remember about areas that have dynamic content). Third, since spatial images of interfaces change over time, we plan to investigate the temporal progression of interfaces images. Fourth, random or excessive use of landmarks in GUIs may have an adverse effect on spatial memory – so we plan to find an optimal number of landmarks to maximize spatial benefits. Last, future studies will compare different landmarks to determine their strengths and weaknesses.

3.3.6 Conclusion

Graphical interfaces display a large number of commands at specific locations through menus and toolbars, but frequent users of these systems can quickly find commands because they already know the locations. In order to figure out how people develop this spatial location memory, we carried out a study with frequent users of four standard GUIs: Word, Facebook, Reader, and Photoshop. Our study revealed that people rely heavily on landmarks readily available in the GUI environment to orient themselves to the GUI, and that these landmarks help frequent users to develop vivid cognitive images of the interfaces. We identify four landmark types that people use in regular GUIs to remember the locations of commands: the layout of an interface provides a high-level reference frame for the objects present in the interface; a clear and unique command grouping is a landmark providing spatial support to recall commands from the group; the corners and edges present in a GUI (both inside and outside) serve as landmarks for nearby commands; and the visual appearance of a command acts as a reliable memory anchor for remembering its location. This work provides new evidence that landmarks can benefit spatial learning and expertise development in graphical interfaces, and provides design guidelines that can help make GUIs more easily memorable.

3.3.7 Acknowledgments

This work was supported by the Natural Sciences and Engineering Research Council of Canada (NSERC). We also thank the anonymous reviewers and committee members for their valuable comments and suggestions.

3.4 SUMMARY OF MANUSCRIPT A

The study presented in this chapter focused on understanding how people develop spatial memory of commands in four standard commercially available graphical interfaces: Microsoft Word, Facebook, Adobe Photoshop and Adobe Reader. The 20-person study consisted of users describing interfaces, drawing sketches, pointing commands and generating walkthroughs to find commands within the interfaces. My findings suggest that users possess vivid images of the interfaces they regularly use. People relied on those cognitive images and the associated landmarks to remember the available items (i.e., commands) and interact with them successfully. Interestingly, no matter how sparse the individual images might appear initially, they seemed to overlap each other, and collectively a public image of an interface was formed.

Results also indicated that the accuracy of the cognitive images varied. For example, the cognitive images of the Word and Facebook interfaces were relatively more accurate than the images of Photoshop and Reader. Surprisingly, Reader had the overall lowest accuracy among the four GUIs, even though it was the simplest interface in terms of the number of commands. One reason could be the linear representation of commands without any clear grouping (a landmark) among the commands. Either way, it is an indication that the lack of adequate landmarks in GUIs makes spatial learning difficult.

Additionally, I have identified that people rely on four different types of landmarks in the standard GUIs, which are available both inside and outside of the visual interfaces. For example, people treated the layout of an interface as a general reference frame (i.e., a landmark) to remember the locations of commands and other elements from the interfaces during the study. A group of commands with a clear boundary served as another powerful landmark in the interfaces. Additional features such as shape, size, colour and meaning of the icons representing commands also supported users in learning and recalling commands' locations. Last, external features, such as edges and corners of screens, including the corners present inside GUIs, served as stable landmarks that users heavily relied on to learn and recall commands' locations during my study.

3.4.1 Contributions

The study presented in this chapter provided three main contributions. First, it revealed for the first time that users who frequently use graphical interfaces could develop spatial images of the interfaces in their minds. This work also demonstrated a novel method to invoke the images of interfaces from users' minds. Besides, my work provided new evidence that the readily available landmarks in GUIs can aid spatial memory development of commands, and their absence can make spatial learning a challenge. Second, I identified four types of landmarks present in the standard GUIs that users heavily relied on to correctly remember the locations of commands. Last, it provided guidelines for designers to make future GUIs more easily memorable by consciously incorporating the idea of landmarks in the interface design process.

3.4.2 Relevance in Context

In the context of my dissertation, Manuscript A has served a greater purpose – it lays the foundation for the research work I report on in the following two manuscripts. A key finding of this manuscript is that it revealed how spatial learning occurs in standard GUIs and identified four landmarks present in these GUIs, which were not considered landmarks before. It was a vital step towards achieving the goals of my dissertation because it solidified our understanding of the spatial memory development process in GUIs and revealed the roles landmarks play in it. It also reconfirmed the problem addressed in this dissertation by demonstrating that inadequate landmarks in GUIs often make it difficult for users to learn and remember the locations of commands; therefore, in Chapters 4 and 5, I investigated the use of artificial landmarks as a novel reference frame in GUIs.

CHAPTER 4

EFFECT OF ARTIFICIAL LANDMARKS ON COMMAND SELECTION

Citation: Md. Sami Uddin, Carl Gutwin, and Andy Cockburn. 2017. The Effects of Artificial Landmarks on Learning and Performance in Spatial-Memory Interfaces. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, ACM, New York, NY, USA, 3843–3855.

Contributions: Under the supervision of Dr. Carl Gutwin, I designed and implemented the prototype command selection interfaces used in the three studies presented in this manuscript, including the logging for the selections during studies. I was also responsible for directing this research, study design, data analyses, and reporting results. Dr. Andy Cockburn was involved in parts of designing studies and preparing the manuscript.

Selecting a command from a menu is a common yet essential task that users regularly perform in the graphical interfaces of computer applications. The goals of Manuscript B were to explore command selection interfaces to understand the effect of artificial landmarks on the performance of command selection – that is, to determine if artificially created landmarks can provide reference frames for spatial memory to improve learning and remembering the locations of commands in graphical interfaces. I carried out three separate yet connected lab studies with prototype interfaces augmented with artificial landmarks and compared them with equivalent non-landmarked interfaces. In these interfaces, I tested the use of landmarks at the level of an entire interface, where landmarks (and commands) were distributed across the whole interface. The findings indicated that artificial landmarks could significantly improve command selection performance in graphical interfaces, particularly when the number of commands increases.

4.1 PROBLEM AND MOTIVATION

Graphical interfaces represent commands as graphic icons and arrange them in menus or toolbars that often place commands in complex hierarchies. Users, particularly novices who are new to a GUI, typically carry out a slow visual search to find the desired command out of many commands available in the GUI. However, experts can quickly locate a command in a GUI by recalling it from memory since they already know its location. One way novices can become experts in a GUI is by learning the locations of commands. However, if an interface fails to support spatial learning, it forces users to rely on slow visual search-based command selection.

Spatially stable menu designs – that display commands at fixed locations – can enable users' quick access to commands by allowing them to anticipate their locations within the interface. Research interfaces, such as Gutwin et al.'s FastTap menu [95], Scarr et al.'s CommandMaps [183], and Gaur et al.'s Multi-Tab FastTap [81], have exploited the benefits offered by the spatial consistency of commands in an interface to improve command selection performance. These interfaces utilize the entirety of a screen to display all commands at once. However, the large number of commands available in GUIs can make it difficult for users, particularly novices, to develop spatial memory of commands partly because GUIs do not provide adequate landmarks.

The literature discussed in Chapter 2 indicated that people learn and remember locations in real-life using landmarks as a reference mechanism. In fact, the development of spatial knowledge begins with the acquisition of landmark knowledge [195]. Moreover, Chapter 3 revealed that users tended to rely on landmarks (though insufficient) to remember the locations of commands in standard GUIs. Therefore, in this chapter, I set out to investigate whether adding landmarks in GUIs can assist users in spatial location learning and recalling. If an application's interface can help users to learn and recall the locations of commands quickly, it can help users improve the efficiency of command selection and contribute to users' expertise development with an interface. In Manuscript B, I focused on learning the locations of commands, so I explored command selection interfaces.

Therefore, the problem addressed in Manuscript B is that graphical command selection interfaces often do not provide enough landmarks to support efficient learning and remembering of commands.

4.2 SOLUTION AND STEPS TO SOLUTION

To address the problem of inadequate referencing mechanisms in command selection interfaces, I investigated two artificially created elements as landmarks in GUIs: coloured blocks in a menu and an image as a menu backdrop. Other elements, such as icons or thumbnails, could also act as landmarks, but I focused in this study on external landmarks. I compared three interfaces (a grid menu with no landmarks, highlighted anchor points within a grid, and an image as the background of a grid) which were designed based on Scarr et al.'s [183] CommandMaps that uses the entire display space of a screen to show all commands at once. I chose this design because it can aid spatial learning and improve command selection performance over standard hierarchical menus [183]. The horizontal and vertical lines of a grid can provide a spatial orientation for commands, particularly in small devices [81,95,132]. However, the repeated use of lines on large screens (e.g., desktops) would reduce the efficiency of the grid as landmarks [222]. So, despite the grid having the potential to be a landmark, the grid menu with no landmark was treated as the control condition among the three interfaces.

The number of commands varies from one interface to another. For example, the document reader application Adobe Acrobat Reader has very few commands in the interface, whereas the document editing application Microsoft Word or photo editing application Adobe Photoshop interface comprises several hundred. To test the use of *artificial landmarks* on commands' locations learning and recalling, I decided to use three different command set sizes: small, medium, and large. Due to the length of the study (three conditions in each of the three sets), it was not possible to test three different command-set sizes in the same study. Therefore, the three sizes were studied in three separate experiments. Study 1 used 64 items in an 8x8 grid. Studies 2 and 3 used 96 commands in an 8x12 grid and 160 items in a 10x16 grid, respectively.

4.2.1 Research Questions

The manuscript consists of three separate studies that tested two landmark types (i.e., anchors and images) in three different sizes of command sets. The studies involved in this manuscript focused on answering the following questions:

- Do artificial landmarks facilitate spatial learning?

- How do artificial landmarks affect the performance of command selection?
- What type of artificial landmarks support better spatial learning?
- Do artificial landmarks perform differently when the number of commands varies?

4.3 MANUSCRIPT B

Spatial memory is a powerful way for users to become expert with an interface, because remembering item locations means that users do not have to carry out slow visual search. Spatial learning in the real world benefits greatly from landmarks in the environment, but user interfaces often provide very few visual landmarks. In this paper we explore the use of artificial landmarks as a way to improve people's spatial memory in spatially-stable grid menus called CommandMaps. We carried out three studies to test the effects of three types of artificial landmarks (standard gridlines, simple anchor marks, and a transparent image) on spatial learning. We found that for small grid menus, the artificial landmarks had little impact on performance, whereas for medium and large grids, the simple anchor marks significantly improved performance. The simple visual anchors were faster and less error-prone than the visually richer transparent image. Our studies show that artificial landmarks can be a valuable addition to spatial interfaces.

4.3.1 Introduction

Spatial consistency is a powerful means for enabling expert performance with user interfaces. By providing stable spatial locations, users can anticipate the location of items and quickly acquire them. Touch-typing is a good example – people can quickly access specific letters without thinking about key locations, and can even form chains of anticipated motor actions that are executed semi-autonomously (e.g., typing the characters of a word while composing the subsequent sentence). When interfaces fail to support spatial consistency, as they often do, users instead need to resort to comparatively slow visual search to find items, negating opportunities for anticipatory action.

Many research and commercial interfaces have been explicitly designed to exploit the efficiencies offered by learned, stable, spatial locations. For example, Marking Menus [130] allow large command vocabularies to be quickly accessed through a fluid series of directional gestures.

CommandMaps [183] also allow access to large command vocabularies, but they do so by flattening the traditional command hierarchy, assigning each command to a unique spatial location in the display. Third, gestural ShapeWriting [233] allows users to input text on mobile devices by sweeping out an approximate gesture over a series of spatially stable characters on a virtual keyboard.

Although spatially stable interfaces can enable high input efficiency once the user is an expert, the attainment of expertise depends on the user learning, remembering, and efficiently recalling item locations. Learning to touch type, for instance, typically consumes months of training, and the skill is refined for years. Few office workers, however, would be willing to engage in such dedicated training to become proficient with a new user interface.

There are therefore important research questions in determining effective methods to assist users in learning and recalling spatial locations. For example, previous studies have examined the learning benefits derived from promoting 'deep encodings' [61] – by removing the continual availability of visual feedback, users are forced to actively engage their spatial memory (rather than rely on visual search) which has been shown to improve users' recollection of the location of abstract icons [72], the location of keys on a new keyboard layout [52], and the shape of command gestures [11,30]. However, these approaches intentionally make interaction harder for novice and intermediate users, which may be acceptable for those with a desire to become expert, but will frustrate many others.

An alternative approach for facilitating spatial learning is motivated by a real world mechanism that novices and experts use to augment spatial memory – landmarks. Landmarks are readily identifiable features in space that are easily discriminated from their surrounds [143] and which serve as an orientation point for spatial actions in a familiar or unfamiliar environment. The use of landmarks has been frequently examined in first-person navigation through 3D virtual environments, such as virtual and augmented reality [64,180]. However, there has been comparatively little research into how landmarking features can facilitate object spatial memory in the static and substantially 2D layouts that dominate mobile and desktop interfaces.

In striving for uncluttered and visually appealing user interfaces, contemporary designs often contain few graphical embellishments. While this may improve aesthetics, it also creates a void of potential landmarks that users might otherwise have employed to assist the formation and use of spatial memory. Other interfaces, in contrast, are heavily populated with features that could be used as landmarks. For example, desktop wallpaper images provide a backdrop that may help users memorize the location of icons. Opportunities for leveraging visual embellishments to assist interaction have been examined in previous work – for example, in observing the disparity between clean user interfaces and messy, dog-eared paper documents, Hill et al. [100] proposed the use of *edit-wear* and *read-wear* to graphically augment interface elements with traces of the user's activity. Hill suggested that these augmentations could indicate frequent and recent activity – but few studies have examined the role that visual information and landmarks can play in assisting spatial interaction with user interfaces.

The three studies reported in this paper examine the influence that different forms of artificially-added landmarks have on spatial learning and recall. Experimental tasks involved retrieving items from a grid-menu of alternatives, similar to CommandMaps [183]. The number of candidate items increased across the three studies (8x8, 8x12, and 10x16). Each study compared item location learning across blocks using three forms of landmarking assistance: a *standard* unadorned grid of icons that used only gridlines for background landmarking; a grid augmented with visual *anchors* in the form of gray backgrounds for a few items in the grid (the intention being to provide clear landmarked reference points for the user); and an *image* background that used a transparent overlay image (of the Taj Mahal) plus background gridlines. The abstract landmarks of the *anchored* condition provide highly distinct spatial demarcation, and the *image* condition offers semantically meaningful features (for example, 'the icon by the right turret') – both of these landmarking aids could potentially offer spatial memory advantages. Our overall hypothesis, then, is that the anchor points and the overlay image will assist learning the locations of items in the grid.

Results showed that *anchor* landmarks were most effective in improving users' spatial memory and performance. There was no difference between the techniques in small grids, but with larger grids, error rates and subjective preferences all favored the *anchor* condition.

The studies provide three main contributions. First, we show that for smaller spatial interfaces, artificial landmarks over and above a basic grid offer little benefit. Second, we demonstrate that as interfaces grow larger, the value of artificial landmarks increases significantly. Third, we provide empirical evidence about spatial learning and spatial retrieval that can assist designers as they build future interfaces based on spatial memory.

4.3.2 Related Work

4.3.2.1 Interfaces for Improved Selection Performance

From entertainment applications to office work, command selection is one of the fundamental tasks that users perform. Selection performance in these interfaces is dependent on two operations. First, users must find a specific command among those available, and second, they must execute that command by pointing to it with a pointing device. Generally, pointing time depends on target width and distance (i.e., Fitts' Law). The time to find a command, however, is related to users' familiarity with the interface [50]. Inexperienced users must rely on slow visual search, but knowledgeable users can skip this step [99] and simply recall the command's location – speeding up performance.

Considerable research has examined methods to improve performance in both of these stages. Alternative command organizations are one main approach: for example, to reduce pointing time, *pie menus* [42] place the commands in a circle around the cursor upon invocation. *Marking menus* [127] use a similar radial organization, but also allow experts to perform pre-emptive gestural selections. Other approaches attempt to flatten command hierarchies to reduce the fixed costs of navigating between levels of the hierarchy (e.g., CommandMaps [183] and FastTap [95] use grid approaches, and other techniques orient items around a user's hand [220]). Keyboard-based shortcuts (i.e., hotkeys [148]) are another way to improve performance [165]; however, studies have shown that real-world use of these tools is often limited [186].

Accommodating a large number of commands within a selection technique is also an important issue, because typical ways of adding commands (e.g., with menus or ribbons) often add hierarchies which slow performance. Memory-based selection should allow a large command set while also maintaining fast access. A few examples exist for high-capacity techniques, such as Marking Menus [127] (64 items or more), ListMaps [92] (225 font items), CommandMaps [183] (210 items), or Kurtenbach et al.'s Hotbox, which supports large command sets by grouping the menu items into different Marking-Menu zones [128]. A problem for all large-capacity memory-based techniques, however, is that remembering command locations may become difficult as command set size increases.

4.3.2.2 Memory-based Interaction

Memory-based interfaces allow users to go directly to a command by recalling its location, rather than by visual-search-based navigation. Human memory is a well-studied topic, both in HCI and psychology (e.g., [63,72,170,208]). Numerous techniques such as gestures [130], hotkeys [148], spatial locations [49,95], or multi-touch chords [82] demonstrated that people can build up extensive mappings between sets of items and command-invocation actions.

Gestures are a popular type of memory-based technique. For example, marking menus [127] and flower menus [29] provide a transition from navigation-based selection to memory-based gestures. In early use of these techniques, items can be found through visual search; but as users repeat selections for common items, they can begin to perform quicker selections by carrying out an accelerated and feedback-free version of the novice method. Other gesture techniques such as Octopocus [30] and Hotbox [128] try to aid the learning.

Past research with spatial memory in computer interfaces has shown that people can remember a large number of locations and can revisit them rapidly. For example, the Data Mountain [177] technique was significantly faster than ordinary bookmarking for retrieval of 100 web pages. Grid-based menus such as ListMap [92], FastTap [95], Square Menus [4], and CommandMaps [183] all showed performance advantages over either search-based or hierarchical organizations of data. However, there is still little understanding of the limits on spatial memory as a basis for user interfaces – and in particular, little understanding of how best to support location learning in these methods as command sets grow larger.

4.3.2.3 Use of Landmarks in Interfaces

In GUI-based systems, the landmarks that are already present in the environment (e.g., the corners of the screen) can provide a strong external reference frame that helps users build up spatial memory [215,217]. Several techniques have explicitly made use of the edges and corners of small devices (e.g., tablets or smartwatches) as landmarks to organize menus and toolbars [95,132,191]. However, these natural landmarks become less useful with larger screens, because many locations are not near landmarks. One technique for tabletops addressed this problem by using real-world objects (the user's own hands and fingers) to provide anchor points for faster location retrieval [220].

When natural landmarks are insufficient, artificially created visual elements can serve as landmarks [17,200]. Artificial landmarks, such as colour [7], can help people to revisit an intended location quickly, and shape has also been used to give an object a memorable "visual ID" [135]. In addition, several techniques add marks to an interface to aid tasks such as understanding activity (e.g., Edit Wear and Read Wear [100]) and revisiting previously-seen items (e.g., Footprints Scrollbar [7] or Visual Popout UIs [78].

Several video summarization systems also use a type of landmark – e.g., creation of storyboards that indicate scene changes. For example, SceneSkim [167] provides browsing and skimming facilities using captions, scripts and plot summaries as reference points for different video locations, and Video Digests [168] represent sections/chapters with navigable markers. Other tools show visual highlights on timelines that represent personal [5] or crowd [124,231] navigation history, to support exploration and revisitation.

The design of landmarks has also been considered in 3D virtual environments to enhance first-person navigation [64,180], and guidelines exist for the design of landmarks for virtual worlds [222]. Less is known, however, about the design or value of landmarks in spatially-stable 2D interfaces. In our work, we use CommandMap's flat menu approach [183] to show a large number of commands in a grid-based overlay menu. We opted for this technique as it provides a basic representation on which people can develop spatial memory, and because it provides ample space for different types of artificial landmarks.

4.3.3 LandMark Test Interfaces

To test the effect of different approaches to landmarking on spatial memory, we designed three similar interfaces (Figure 4.1) based on CommandMaps [183]. CommandMaps use all of the available display space to concurrently reveal all of the commands, each shown in a unique and stable spatial location. Normally, the CommandMap is not shown, allowing the full display space to be dedicated to the user's workspace (such as a document or spreadsheet). However, when the user wishes to access a command, they issue a control command (such as pressing a modifier key, mouse button, or gesture), which causes the CommandMap to be revealed. Command items are then selected by pointing and clicking on them. The CommandMap can be hidden either after each command selection, or on a subsequent control action (possibly allowing multiple commands to be invoked in a series, if the application requires it). As users learn the location of items in the CommandMap, they can anticipate the location at which they will be presented, facilitating rapid selection. Previous lab studies have demonstrated the efficiency of CommandMaps in comparison to menus and toolbar interfaces, both in abstract and realistic tasks [184].

Figure 4.1: Interfaces with artificial landmarks (Study 1): (left) Standard menu with grid, (middle) Anchor menu with anchor points, and (right) Image menu with Taj Mahal's image as background.

As shown in Figure 4.1, our three landmark interfaces initially show a grid (*standard*, left), a grid plus a small set of dark gray grid anchors (*anchor*, middle), or a grid plus a transparent image of the Taj Mahal (*image*, right). When the user presses the Control key, a set of underlying icons are revealed in full screen setup. Selections are then made by clicking on the appropriate icon. Icons remain displayed until the Control key is released, or an icon item is clicked.

All three interfaces also support an *expert mode* of selection, in which icons can be selected prior to their display by pressing the Control key and immediately clicking in the location corresponding to the target item. To facilitate and encourage expert selections all interfaces implemented a timeout (200-400ms depending on the command set size) between pressing the Control key and displaying the icons.

We use the terms '*basic mode*' for selections that are completed with the aid of visual feedback after the short timeout, and '*expert mode*' for selections completed prior to the display of icons (Figure 4.2). As users' spatial memories of icon locations improve, they should complete more selections in *expert mode*. The landmarks used in the interfaces were always available in both *basic* and *expert mode* of selections (as shown in Figure 4.2).

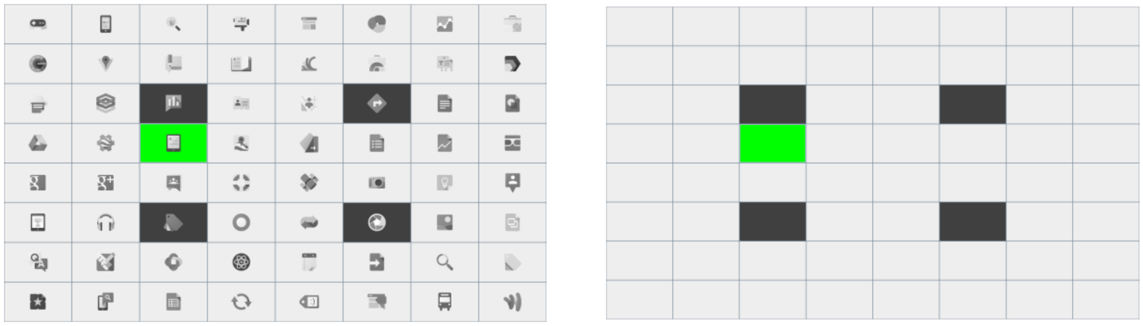

Figure 4.2: Selection modes: (left) basic - selection after icons are shown, (right) expert - selection without seeing the icons.

4.3.3.1 Standard Grid

The standard grid provides clear borders that demark item locations. In grids with relatively few items it is likely that the coarse placement resolution will allow the borders and corners to provide sufficient inherent landmark cues to assist spatial memorization (e.g., 'item by the top right corner').

4.3.3.2 Anchor

As the number of grid items increases, it is likely that the inherent landmark cues provided by borders and corners will become less effective in aiding spatial memorization and retrieval. Similarly, the grid lines provided by the *standard* condition are also likely to become less effective due to the frequency of their repetition.

The *anchor* condition therefore augments the spatial grid with dark grey grid cells that provide clear spatial reference points and are distinct from their surroundings [222]. Anchor cells are visible both before and after the icons are displayed. They therefore provide strong spatial anchors in *basic mode* selections, and particularly in *expert mode* where selections can be made without waiting for the icons to appear.

4.3.3.3 *Image*

The abstract landmarks provided in the *anchor* condition may provide inferior support for memory formation when compared to semantically meaningful objects. To test this possibility, the *image* condition overlays the grid with a simple greyscale image of a building (the Taj Mahal) with clear features (such as turrets, windows, doorways, paths, etc.). These features may assist memorization through the opportunities for an association between objects (the icons) and the meaningful spatial location in the image (in a manner similar to the 'method of Loci' [169,230,245]). The inclusion of both an artificial landmark and image condition is partially motivated by prior findings on contextual cueing, which have indicated differences between item location learning in naturalistic scenes versus simple stimulus arrays [38,39].

4.3.4 Study 1: Landmarks in a Small Command set

The following three studies examined user performance with our three landmark interfaces using progressively larger grids and with different lengths of training to enable memorization. Study 1 used an 8x8 grid of 64 items, and 11 blocks of trials. (Study 2 used a 8x12 grid (96 items) and 16 blocks of trials; Study 3 used a 10x16 grid and 18 blocks). All three studies were designed to answer the same main question: do the different approaches to landmarking result in different performance and styles of use?

4.3.4.1 *Study 1 – Method*

Tasks and stimulus

The study consisted of a series of trials, each involving the 'point and click' selection of a cued icon stimulus. Each trial began by displaying the stimulus icon on the left screen of a dual-monitor (21-inch) environment. The participant then used a mouse to select the target icon using one of the three landmarking interfaces (Figure 4.1) described above (*standard, anchor,* or *image*) which ran full-screen on the right screen.

The 64 icons (64px in size) presented in the interfaces were the same across all three interfaces (randomly relocated for each). They were extracted from the Android icon set [246] and converted to grayscale to reduce potential confounds from hue-induced popout effects. Twelve of the icons were quasi-randomly selected for use as stimuli with each interface condition. None of the icon stimuli were reused with a subsequent interface condition.

Procedure and study design

Participants were initially informed that the experiment concerned interfaces for rapid command selection. The *basic* and *expert* selection modes were described and demonstrated. They then completed 20 practice selections with each of the three interfaces using a different dataset to that used in the main experiment. Participants were instructed to complete trials as quickly and accurately as possible.

Participants completed 11 blocks of trials with each of the three interface conditions (order was counterbalanced). Having completed all blocks with one interface, participants completed a NASA-TLX [96] subjective workload questionnaire. They then progressed to the next interface.

Each of the first 10 blocks consisted of one trial for each of the 12 targets. Users were free to complete each trial using whichever selection modality they preferred (*basic* or *expert*). The 11^{th} block involved a ‘blind’ trial for each target, with the *basic* selection mode disabled – participants clicked on the location they believed corresponded with the cued icon, without visual feedback.

Within each trial, a correct selection was confirmed by highlighting the selected grid location green for 400 ms; red for incorrect. Trials continued until correctly completed. Software recorded trial completion time, errors, expert selections, and data describing every selection. At the end of the study, participants chose their preferences for the three interfaces for various aspects of interaction.

Participants

Twelve participants (1 female), ages 18-34 (mean 25.5), were recruited from a local university. The study took ~60 minutes, and a $10 remuneration was paid to each participant.

Apparatus

The experiment was conducted on a desktop computer running Windows 7, with two 21-inch 1650x1050 resolution monitors placed alongside. Software was written in Java. Input was received through a standard keyboard and optical mouse. All study interfaces ran full-screen in the right of two 21-inch monitors.

4.3.5 Study 1 – Results

For all the three studies, we report the effect size for significant RM-ANOVA results as partial eta-squared: η^2 (considering .01 small, .06 medium, and >.14 large [30]). In all studies, where ANOVAs sphericity assumption is violated (Mauchley's test), Greenhouse-Geisser adjustments are performed (yielding floating point degrees of freedom).

4.3.5.1 Trial time

Mean trial completion times with the three interfaces are summarized across blocks in Figure 4.3.

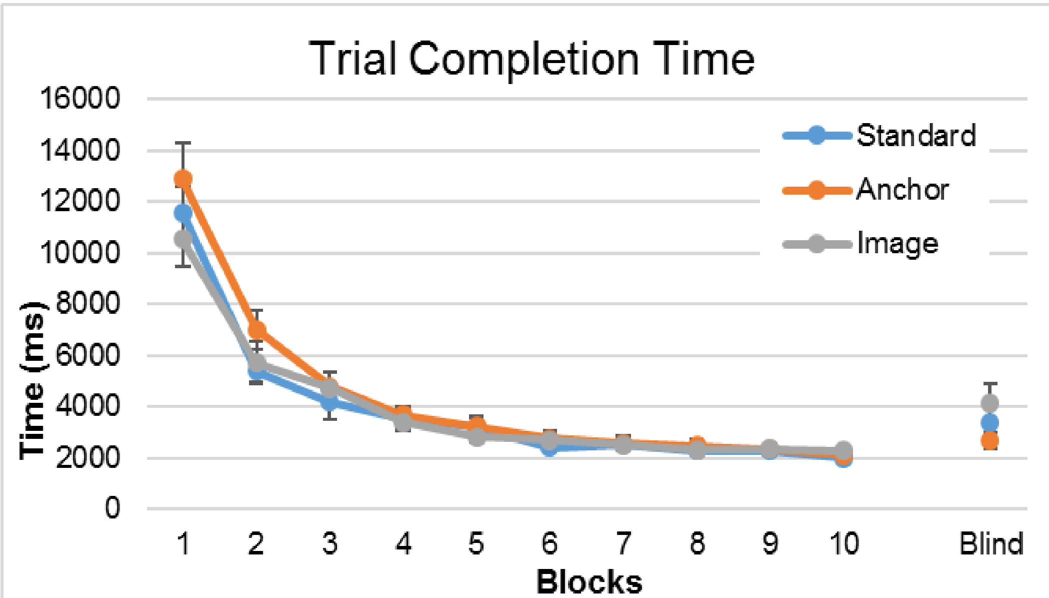

Figure 4.3: Mean trial time (±s.e.) by interface and block.

For the ten main blocks, RM-ANOVA showed no significant main effect of *interface* ($F_{2,22}$=2.0, p=.16), with means of 3940ms (s.d. 3146ms) with *standard*, 4348ms (s.d. 3727ms) with *anchor*, and 3931ms (s.d. 3003ms) with *image.*

Trial times decreased across *block* ($F_{1.84,20.2}$=78.84, p<.001, η^2=0.88), and as anticipated, the skill development follows a power-law function [159]. There was no significant *interface* × *block* interaction ($F_{18,198}$=1.04, p=.42).

Analysis of mean trial time in the final blind block also showed no significant difference between the three interfaces ($F_{2,22}$=2.1, p=.15).

4.3.5.2 Error rates and expert selections

Analysis of the number of errors per trial showed similar results to the trial time analysis. There was no main effect of *interface* ($F_{2,22} < 1$) and no *interface* × *block* interaction ($F_{18,198} < 1$). Errors increased with *block* ($F_{9,99}$=2.94, p=.004, η^2=0.21), which can be attributed to two potential causes – users becoming faster and less precise, and users increasingly attempting to rely on incompletely-formed spatial memories with the expert mode.

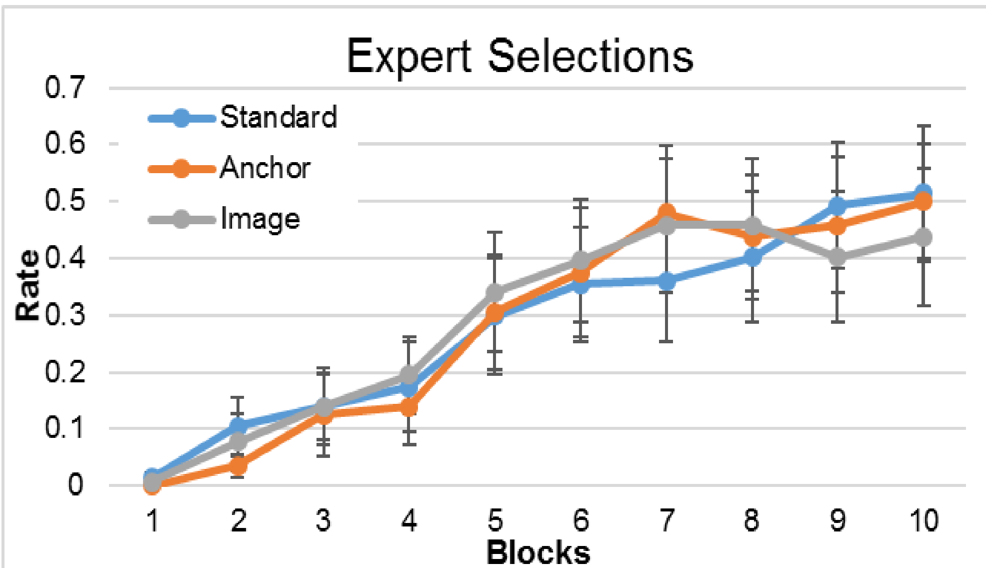

Figure 4.4: Expert selection rates by interface and block.

Use of the expert selection mode followed a similar pattern to errors, with only *block* ($F_{1.63,17.92}$ =14.96, p<.001, η^2=0.58) showing a significant effect (Figure 4.4). Participants made the same proportion of expert selections with all three interfaces (overall, 0.29 selections/trial).

4.3.5.3 Subjective responses

Participant responses on the NASA-TLX worksheets were also similar for the three interfaces (see Table 4.1), with no significant effects except for reported Frustration: *image* (mean 3.25, s.d. 2.3) induced higher frustration than *anchor* (2.8, s.d. 2.7) and *standard* (2.3, s.d. 1.2).

Table 4.1: Mean (s.d.) effort scores (0-10 scale, low to high).

	Standard	Anchor	Image	χ_r^2	p
Mental	7.17(1.67)	6.58(1.8)	6.42(1.38)	1.79	.41
Physical	2.67(2.49)	2.67(2.32)	3.42(2.6)	1.54	.46
Temporal	5.25(2.05)	5.75(1.69)	5.5(2.06)	0.79	.67
Performance	5.00(2.74)	5.33(2.95)	4.25(2.45)	3.16	.21
Effort	6.00(2.31)	6.00(1.91)	5.92(1.61)	0.13	.94
Frustration	2.25(1.23)	2.75(2.68)	3.25(2.28)	3.12	**.01**

Table 4.2: Count of participant preferences.

	Standard	**Anchor**	**Image**
Speed	1	8	3
Accuracy	1	8	3
Memorization	0	8	4
Expertise	3	7	2
Comfort	1	9	2
Overall	1	8	3

Despite the lack of objective and subjective workload findings favouring any interface, the participants' preferences were strongly in favour of the *anchor* interface. Eight or more of the twelve participants selected it as the preferred interface for Speed, Accuracy, Memorization, Comfort and Overall (see Table 4.2).

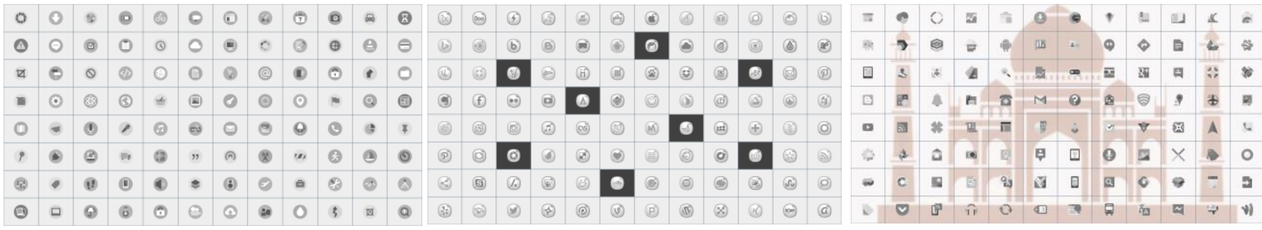

Figure 4.5: Study interfaces for Study 2: (left) Standard menu with grid, (middle) Anchor menu with anchor points, and (right) Image menu with Taj Mahal's image as background.

4.3.6 Study 2: Landmarks in a Medium Command Set

The second study used a similar method to Study 1, but with the following alterations. First, the grid used in all three interfaces was larger, with 96 items (64px) in a 8x12 arrangement; the item set was unique in each condition to avoid learning effects. Second, there were 15 blocks of trials instead of 10; and a final 16th 'blind' block was used in the same manner as Study 1. Third, nine targets were used in each block rather than 12; this adjustment, together with the increase in the number of blocks allows greater opportunity for spatial learning with all interfaces. Fourth, in the *anchor* condition there were 8 dark gray items rather than 4, positioned as shown in Figure 4.5. Last, we increased the timeout delay to 350ms from 200ms. These adjustments were made to accommodate the higher number of items.

In other aspects, the method, procedure, and apparatus were identical to Study 1.

Participants

Twelve participants (5 females), aged 19-37 (mean 27), were recruited at a local university. None had previously participated in Study 1. The study lasted for approximately 60 minutes, and a $10 honorarium was paid to each participant.

4.3.7 Study 2 – Results

4.3.7.1 Trial time

Mean trial times with the three interfaces across blocks are shown in Figure 4.6. RM-ANOVA showed no significant main effect of *interface* ($F_{2,22}$=3.38, p=.05). Mean trial times were lowest with *anchor* (3990ms, s.d. 4643ms), followed by *standard* (5057ms, s.d. 4464ms) and *image* (7463ms, s.d. 4204ms).

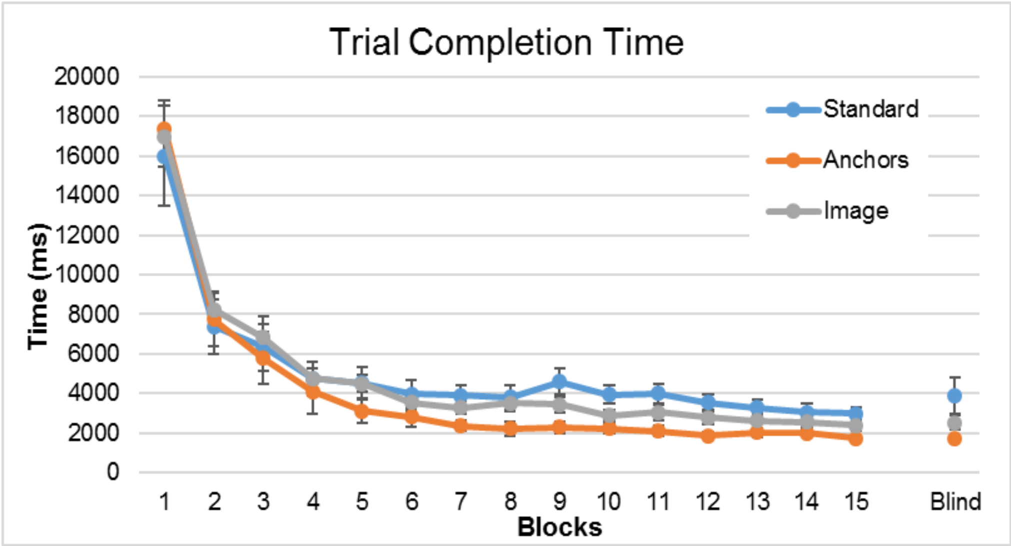

Figure 4.6: Mean trial completion time by interface and block.

As in Study 1, there was a significant effect of *block* ($F_{2.35,25.87}$=64.1, p<.001, η^2=0.85), but no *interface* × *block* interaction ($F_{28,308}$=1.1, p=.34).

Unlike Study 1, analysis of mean trial time in the final *blind* block showed a significant difference between the three *interfaces* ($F_{1.13,12.45}$=3.75, p=.07, η^2=0.25), with *anchor* being the fastest (mean 1717ms, s.d. 309ms).

4.3.7.2 Error rates

There was a significant main effect of *interface* on errors ($F_{2,22}$=9.2, p=.001, η^2=0.46). *Anchor* had the lowest error rate at 0.07 errors/trial (s.d. 0.13), compared to much higher error rates of 0.23 (0.25) and 0.27 (0.3) errors/trial with *standard* and *image* respectively. Bonferroni-corrected follow-up t-tests showed that *anchor* was more accurate than both of the other interfaces (all p<.001).

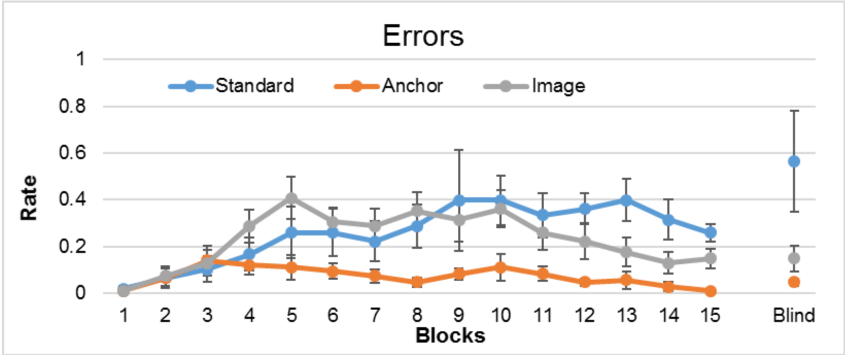

Figure 4.7: Error rate across blocks with the three interfaces.

As shown in Figure 4.7, with all of the interfaces, errors increased across the first blocks, then roughly stabilized, before decreasing in the final blocks, leading to a significant effect of *block* ($F_{4.09,44.97}$=5.99, p<.001, η^2=0.35). There was a significant *interface* × *block* interaction ($F_{28,308}$=2.17, p=.001, η^2=0.16), attributable to *Anchor* having comparatively stable and low errors across blocks.

Table 4.3: Percentage of errors at each target location group (side, corner or middle), broken down by error distance.

	Standard		Anchor		Image	
Group	**Off by 1**	**Off by 1+**	**Off by 1**	**Off by 1+**	**Off by 1**	**Off by 1+**
Side	14.4	3.6	3.9	1.9	12.8	5.8
Corner	7.5	3.9	2.8	1.9	5.0	1.7
Middle	23.7	10.6	2.1	6.6	21.7	9.8

We also analyzed error rates based on the location of target items in the grid (along the side, at the corner, or in the middle), and we categorized errors based on the distance from the intended target (‘off by 1’ and ‘off by more than 1’). Table 4.3 summarizes the findings, showing the proportion of trials at each location that contained an error at each distance. The table reveals some interesting additional characteristics of errors. With *standard* and *image* the highest error rates occurred in the middle of the grid. With *standard*, the total error rate for middle targets was 34.3% (23.7+10.6), and with *image* it was 31.5%. These high rates contrast with the relatively low value of 8.7% for middle targets with *anchors*. Corner error rates were much lower with all three interfaces (*standard* 11.4, *anchors* 4.7, and *image* 6.7%), which can be explained by the unambiguous spatial demarcation provided by the corner. Interestingly, side errors were also much lower with *anchors* (5.8%) than *standard* (18.0%) and *image* (18.6%), which may be due to the *anchors* providing clear reference points along both independent dimensions (e.g., "on the right edge and aligned with the top gray block").

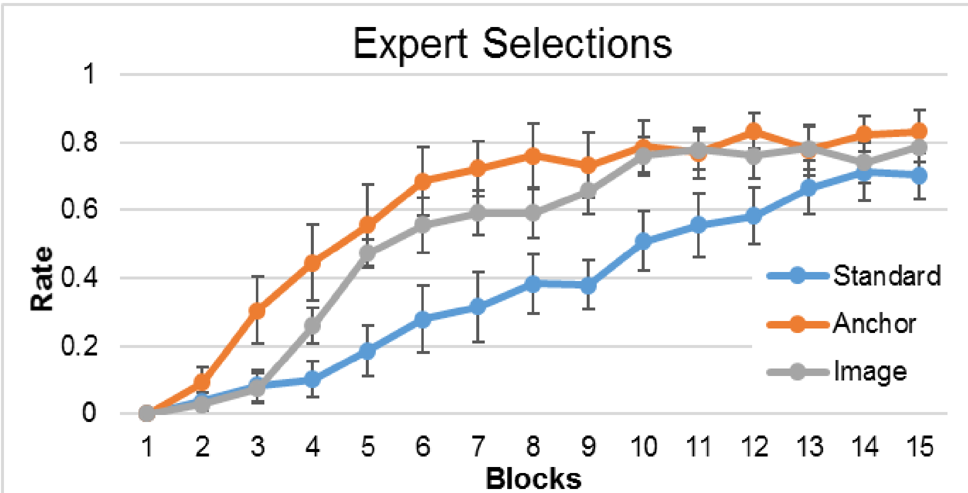

Figure 4.8: Expert selection rates by interface and block.

4.3.7.3 *Use of expert selections*

Figure 4.8 shows the rate of expert mode selections with the three interfaces across block. It shows that the *standard* interface had much lower use of the expert mode than the other interfaces in nearly all blocks (mean 0.37 sel./trial, s.d. 0.35), and that the *anchor* interface had the highest rate (mean 0.61 selections/trial, s.d. 0.38); expert mode selections with *image* were slightly lower (0.52 sel./trial, s.d. 0.34): $F_{2,22}$=6.74, p=.005, η^2=0.38). There was a significant *interface* × *block* interaction ($F_{28,308}$=2.27, p<.001, η^2=0.17), as indicated in Figure 4.8. Post-hoc t-tests (Bonferroni corrected) show that *anchor* performed significantly better than the other two interfaces, and that *image* was also faster than *standard* (all p<.001).

4.3.7.4 Subjective responses

Table 4.4 summarizes mean response to the NASA-TLX worksheets. Friedman tests showed significant effects for Mental workload (with *image* lowest, followed by *anchor* and *standard*), Temporal workload (*anchor* lowest, then *image* and *standard* highest) and Performance (*anchor* highest, *standard* lowest).

Table 4.4: Mean (s.d.) effort scores (0-10 scale, low to high).

	Standard	**Anchor**	**Image**	χ_r^2	***p***
Mental	7.17(2.43)	5.42(2.33)	4.33(2.17)	9.04	**.01**
Physical	3.17(3.21)	2.92(2.93)	2.75(2.77)	0.29	.86
Temporal	5.42(2.75)	4.00(2.86)	5.08(2.72)	6.5	**.04**
Performance	5.75(1.69)	7.25(1.3)	6.00(1.96)	10.79	**.01**
Effort	6.83(1.91)	5.33(2.87)	6.42(2.14)	5.79	.06
Frustration	4.00(2.58)	2.75(2.42)	3.33(3.06)	3.79	.15

Counts of the preferred interface strongly favoured the *anchor* interface, with 83-92% of participants selecting it as preferred across six dimensions (see Table 4.5).

Table 4.5: Count of participant preferences.

	Standard	**Anchor**	**Image**	**None**
Speed	0	10	2	0
Accuracy	0	10	2	0
Memorization	0	11	1	0
Expertise	0	10	1	1
Comfort	0	11	1	0
Overall	0	11	1	0

4.3.8 Study 3: Landmarks in a Large Command Set

The third study used the method of Study 2, but with a grid of 160 items (48px) in a 10x16 grid arrangement. Additionally, to provide stronger insights into the formation of spatial memory, we used a blind block of trials after every 5^{th} block. There were therefore 18 blocks in total, consisting of three repetitions of 5 regular blocks followed by one blind block. Ten targets were used instead of the nine in Study 2. Also, we increased the number of gray blocks in the *anchor* condition from 8 to 14 (positions are shown in Figure 4.9). Finally, we set the timeout delay for expert mode selection to 400ms. The method, procedure and apparatus were identical to Study 2 in other aspects.

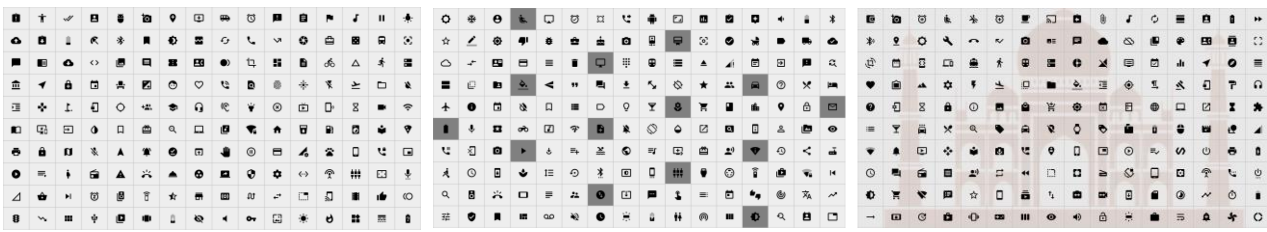

Figure 4.9: Study interfaces for Study 3: (left) Standard menu with grid, (middle) Anchor menu with anchor points, and (right) Image menu with Taj Mahal's image as background.

Participants

We recruited 16 new participants (9 females), ages 19-41 (mean 28.4), from a local university. The study lasted ~60 minutes. Participants received a $10 payment.

4.3.9 Study 3 – Results

4.3.9.1 Trial time

Mean trial times for the 15 *basic* blocks and 3 'blind' blocks (analyzed separately) are shown in Figure 4.10. RM-ANOVA showed a significant main effect of *interface* on completion time ($F_{2,30}$=4.4, p=.02, η^2=0.23). In the *basic* blocks, *anchor* was the fastest with mean 5108ms (s.d. 5620ms) compared to *image* (5215ms, s.d. 4515ms) and *standard* (6416ms, s.d. 6196ms). Bonferroni-corrected t-tests showed that both *anchor* and *image* were significantly faster than *standard* (p<.001), but there was no difference between *anchor* and *image* (p=.68).

Trial time significantly decreased across *block* ($F_{2.1,31.5}$=66.8, p<.001, η^2=0.82), and there was also a significant *interface* × *block* interaction ($F_{28,420}$=2.5, p<.001, η^2=0.14). For the three *blind* blocks (see Figure 4.10), there was no main effect of *interface*, ($F_{2,30}$=2.5, p=.1), but there was a significant effect of *block* ($F_{1.19,17.91}$=9.98, p=.004, η^2=0.4), with mean times becoming slightly faster in later *blind* blocks. There was no *block* × *interface* interaction.

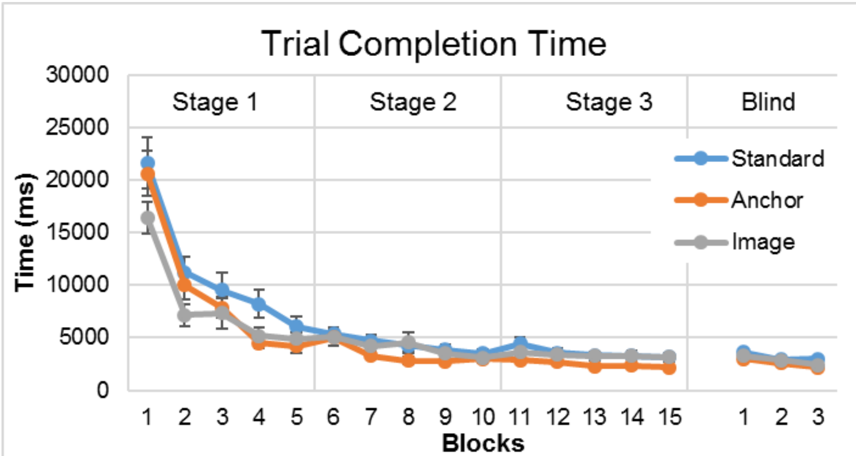

Figure 4.10: Mean trial completion time by interface and block.

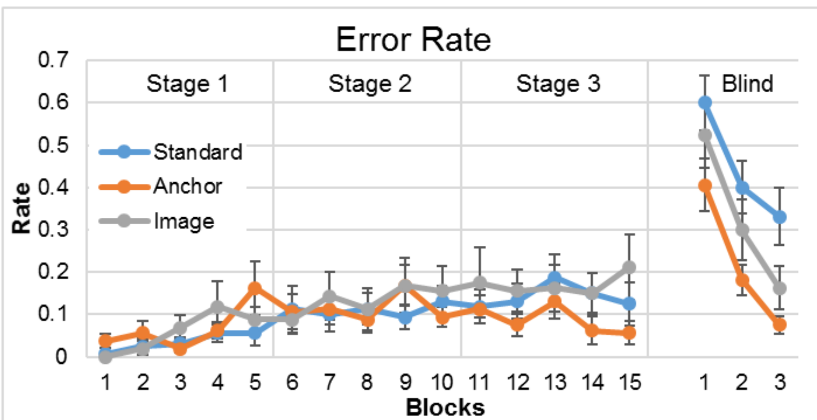

Figure 4.11: Mean error rates in Study 3.

4.3.9.2 Error rates

For the 15 *basic* blocks, there was no significant main effect of *interface* (see Figure 4.11): $F_{1.37,20.52}$=1.52, p=.24. The lack of a significant effect of *interface* differs from Study 2, and is probably due to the additional difficulty of learning 10 items among 160 (rather than 9 among 96), causing more random variation in Study 3. There was a significant effect of *block* ($F_{4.56,68.46}$=3.53, p=.008, η^2=0.19,), but no *interface* × *block* interaction ($F_{28,420}$=0.84, p=.71).

For the *blind* blocks, *anchor* had the lowest error rate (0.22 errors/trial, s.d. 0.22), followed by *image* (0.33, s.d. 0.31) and standard (0.44, s.d. 0.28), giving a significant effect of *interface* ($F_{2,30}$=9.88, p=.001, η^2=0.4). Errors significantly decreased across the three blocks ($F_{2.1,31.5}$=36.18, p<.001, η^2=0.71), with participants making far fewer errors as they learned spatial locations. There was no *interface* × *block* interaction.

Analysis of the influence of target locations (side, corner and middle) on the proportion and distance of errors revealed similar observations to those reported in Study 2 (see Table 4.6). For all interfaces, the highest proportion of errors occurred with *middle* targets. With the *anchor* interface, *middle* errors were lower (14.6%) than *standard* (19.4%) and *image* (20.2%) interfaces. *Corner* errors were similar across interfaces (4.3, 6.7, and 6.4 with *standard*, *anchor*, and *image*). *Side* errors were also lower with *anchors* (5.2%) than with *standard* (14.4%) and *image* (10.9%).

Table 4.6: Percentage of errors at each target location group (side, corner or middle), broken down by error distance.

	Standard		**Anchor**		**Image**	
Group	Off by 1	Off by 1+	Off by 1	Off by 1+	Off by 1	Off by 1+
Side	9.9	4.5	3.5	1.7	6.6	4.3
Corner	4.0	0.3	3.1	3.6	5.4	1.0
Middle	13.1	6.3	3.1	11.5	10.5	9.7

4.3.9.3 Use of expert selections

As shown in Figure 4.12, during *basic* blocks, *anchor* had the highest level of expert selections (mean 0.44 sel./trial, s.d. 0.39) compared to *image* (0.39, s.d. 0.4), and *standard* (0.26, s.d. 0.34): $F_{2,30}$=7.39, p=.002, η^2=0.33. In the final blocks of Stage 3 (when users had had the greatest opportunity to learn locations), approximately 80% of *anchor* selections were made using the expert modality, compared to 65% with *image* and 50% with *standard*. Post-hoc t-tests (Bonferroni corrected) showed differences in expert selection rate for all *interface* pairs (all p<.001).

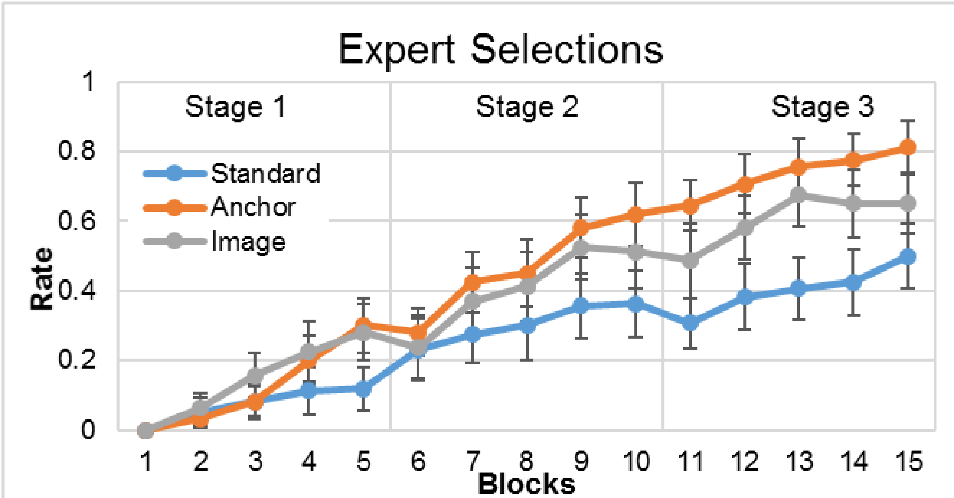

Figure 4.12: Expert selection rates by interface and block.

4.3.9.4 Subjective responses

NASA-TLX responses were also similar to Study 2 (see Table 4.7). *Anchor* and *image* interfaces received lower workload scores and higher performance scores. Preference counts also favoured the *anchor* interface, with no participants selecting *standard* as preferred (see Table 4.8).

Table 4.7: Mean (s.d.) effort scores (0-10 scale, low to high).

	Standard	**Anchor**	**Image**	χ_r^2	***p***
Mental	7.38(1.73)	6.38(2.06)	6.13(2.18)	7.63	**.02**
Physical	2.69(2.49)	2.5(2.06)	2.31(2.02)	0.38	.83
Temporal	5.75(2.68)	4.94(2.49)	4.56(2.6)	6.84	**.03**
Performance	4.44(1.73)	7.19(2.01)	6.75(1.92)	14.28	**<.01**
Effort	7.00(2.18)	6.00(1.9)	6.06(1.52)	5.34	.07
Frustration	5.94(2.9)	4.81(2.74)	4.94(2.84)	4.72	.09

Table 4.8: Count of participant preferences.

	Standard	**Anchor**	**Image**	**None**
Speed	0	8	7	1
Accuracy	0	8	6	2
Memorization	0	8	7	1
Expertise	0	9	4	3
Comfort	0	8	6	2
Overall	0	8	6	2

4.3.10 Participant Comments

Participant comments for the three studies mirrored and emphasized the objective findings. Participants made several comments on how the artificially planted landmarks (especially *anchors*) helped them to develop spatial memory of the commands and improve performance: one participant mentioned *"[Target] locations were easy to predict because of the anchors [in the grid]."* Another said *"It was easy to remember 5-7 anchors and it helped me to find other [target] items."* One person, however, remarked on the difficulty of remembering commands in the *standard* condition: *"Too many items [in grid with no landmark] made it hard to remember the position, although I tried to use grid number to remember [targets]."*

Other comments suggested that the *image* also helped participants to learn command locations: one said *"It was easier to remember things when there was meaningful content [in image] to connect it to (e.g. reading a book on the second floor)."* Another observed that *"I associated the objects with various parts of the monument [image], as if the library was located in one of the minarets of the Taj Mahal."*

However, landmarks can cause distraction and often require extra memory to process information. One person stated *"[Often] when I started looking for any icon, [the] background image caught my attention."*

4.3.11 Discussion

The main findings of the three studies are as follows:

- *Anchor* landmarks improved users' ability to memorize grid item locations, compared to an unadorned grid. By providing abstract spatial cues, users were able to more quickly and accurately select target items.
- In a relatively small grid of 64 items, there were no significant differences between performance with the *standard*, *anchor*, and *image* interfaces. However, preferences strongly favoured the *anchor* interface. It appears that inherent spatial cues provided by the *standard* interface (including the grid, display edge and corners) were sufficient to support effective memorization.
- However, in larger grids, the additional landmarking features provided by the *anchor* and *image* interfaces enabled better memorization and retrieval. Analysis of the proportion of errors by target location (middle, edge, or corner of the display) indicated that errors with the *standard* interface were most prevalent when spatial landmarks were less clear – errors were lowest at the spatially unambiguous display corners, higher along the edges, and highest in the middle of the display (away from corners and edges).
- Mean selection times and error rates were lower with the abstract spatial landmarks provided by *anchor* than they were with *image.* This is interesting because we initially suspected that the *image* interface might present more opportunities for semantic association (e.g., 'the star is by the left minaret), but this was not the case. In addition, *anchor* was strongly preferred over *image*.

4.3.11.1 All Interfaces Supported Spatial Learning

The reduction in mean trial completion time across blocks conforms to the expected power-law of learning curves [159] in all three studies. This can be attributed to users transitioning from selections that are dominated by slow visual search (followed by rapid pointing) to selections that are characterized by rapid spatial recollection and pointing. This transition occurs with all three interfaces, indicating that users developed spatial memories, regardless of the interface. However, Figures 4.8 and 4.12 from studies 2 and 3 show that participants were able to form and exploit these memories more rapidly when additional landmarking features were available in *anchor* and *image* interfaces.

4.3.11.2 Why did Anchor Outperform Image?

Trial time data, error rates, workload measures and preferences all favoured the *anchor* interface over *image.* Yet centuries of evidence from the 'method of loci' [169,245] suggest that concrete spatial representations (such as buildings) can assist the memorization of abstract concepts by associating those concepts with the spatial representation.

We see three main reasons for the comparatively strong performance of the *anchor* interface compared to *image*. First, like many user interfaces, the items used in our study were presented aligned to a clear two dimensional grid. The gray blocks of the *anchor* condition therefore provided strong alignment cues on each dimension – even when a target was distant from one of the blocks, it could serve as a spatial cue ("same row as that gray block over there"). In contrast, the Taj Mahal image less clearly afforded this form of dimension-based alignment – participants may have been less likely to exploit independent row and column alignment concepts when targets were distant from image features.

Second, the gray blocks of the *anchor* interface had high visual salience – users were almost certain to incorporate these landmarks into their conception of the tasks (e.g., "the star is the black top-left block," or "the star is the item above the black top-left block"). In contrast, the Taj Mahal image was more subtly presented, and it might have been ignored by participants during their tasks. Indeed, one participant commented *"I did not notice the background image in grid."* In addition, the finer-grained details on the image may have been more difficult for users to incorporate in their memory, since there was more to remember about the landmark (e.g., "near to this particular filigree on the roofline of the building").

The *anchor* and *image* conditions can be viewed as representing different points on two continuums between sparse and dense image features, and between abstract and concrete representations. Further study is needed to examine the role of these visual characteristics on spatial memorization in user interfaces – but our study suggests that there could be value in simpler representations when providing landmarks that are adjunct to the primary task.

Third, although the Taj Mahal is a well-known building, none of our participants had personal experience of it. The method of loci, however, is based on the intentional placement of memorization objects into a highly personal environment (such as a favourite walk or the rooms of one's home). Results might differ with a more personally-familiar image.

4.3.11.3 Implications for Design

The findings suggest that users' ability to form and draw on spatial memories for rapid interaction will be assisted by the presentation of landmarks. Many interface components that rely on spatial interaction are often featureless, blank spaces, creating problems for users in exploiting their spatial memories. One example is the blank trough of a scrollbar, which can give rise to inefficient interactions. Previous researchers have proposed augmenting the scroll-trough with transient markers to indicate recently or frequently visited regions, which assist users returning to their place in a document after cross-referral to another area [7,100]. Our results suggest that static embellishments in the scroll-trough could provide useful landmarks for associating spatial locations.

Of course, there are challenges to the use of landmarks in user interfaces. The first of these involves the potential for interference from the artificial landmarks (which are always present on the screen) and the user's document content. If the user is working in a graph editor that shows rectangular blocks on the screen, for example, the anchor visualizations could conflict with task objects. This could result in interference in both directions: the artificial landmarks could hinder the interpretation of the document contents, and document graphics could potentially interfere with the value of the anchor landmarks. We believe that this issue can be readily addressed through careful design of the artificial landmarks. For example, the landmarks can be extremely faint (e.g., using a very high level of transparency) and still be useful – once the user is familiar with them, they will need only a minimal representation to guide their spatial memory. In addition, the visual representation of anchor landmarks can potentially be changed without affecting the user's memory – for example, specific colours or textures could be chosen so that they do not interfere with the visual features of the document content.

A related challenge in the use of artificial landmarks involves aesthetics – that is, there is a potential impairment of design aesthetics due to the presentation of otherwise superfluous visual features, and the risk that users will interpret the landmarks as being part of the design of the application rather than as spatial reference points. Further work, particularly in collaboration with graphic designers, is needed to address these concerns.

Although *image* provided richer landmark features than *anchor*, surprisingly, it did not perform as we expected. The Taj Mahal images we used in our studies were carefully converted into grayscale and later faded out to avoid colour overlap with any task icons. It is possible, however, that the background image may have still interfered with a few of the icons – we will consider this issue in further item-by-item analyses. However, *image* provided an overall performance advantage beyond that attained with the *standard* grids. It will be an interesting to investigate the effect of different feature rich images as artificial landmarks.

In future work, we plan to carry out further studies of the value of artificial landmarks, in practical settings with realistic document content, and in different types of applications. For example, the ideas presented here can be tested in one-dimensional representations such as scroll bars and video timelines, as well as in two-dimensional settings such as desktop/homescreen wallpaper decorations. We will also carry out new studies to explore potential interference between artificial landmarks and other objects on the screen, and we will test new designs to determine whether subtle landmark representations that are visually unobtrusive can still provide effective anchors for spatial memory.

4.3.12 Conclusion

Desktop and mobile user interfaces make heavy use of spatial organization to facilitate rapid access to interface items. We examined the role that landmarks play in assisting spatial memorization and retrieval of items in a grid of interface components. The landmarks were static and passive visual embellishments designed to help users orient themselves in the graphical layout. Three forms of landmarking assistance were empirically compared – basic gridlines, the additional use of grey fill for some grid cells to provide clear visual anchors, and the grid overlaid with a meaningful background image. Item retrieval times and error rates were best when using the simple visual anchors. Reasons for the findings, implications for design, and directions for further work were presented.

4.3.13 Acknowledgments

This work was supported by Natural Sciences and Engineering Research Council of Canada.

4.4 SUMMARY OF MANUSCRIPT B

I carried out three connected studies where command set size progressively increased, and two artificial landmarks were tested. The results from the studies revealed that the value of artificial landmarks in supporting location learning and recalling is proportional to the command set size. In the small interface study, consisting of 64 commands, no performance differences were found for any of the interfaces. It indicates that the Standard interface's inherent spatial cues, such as the grid lines, interface's edges and corners, were sufficient to provide stable reference frames for the spatial memory that supported efficient learning and recalling of command locations when command set size was small. In larger command sets (Study 2 and Study 3), however, both artificial landmarks (i.e., Anchors and Images) contributed to better memorization and recall of command locations. A surprising finding was that the Anchor conditions' simple landmarks outperformed the benefits of visually rich image landmarks in developing memory of commands. Participants also made fewer errors while selecting commands in landmark augmented interfaces, particularly with simple anchor landmarks. They also preferred landmarked interfaces over the GUI with no landmarks.

Overall, my analyses indicated that artificial landmarks can be a valuable addition to command selection interfaces, especially when the command set size increases, and that landmarks can support spatial memory development of commands.

4.4.1 Contributions

The main contribution of this work provided empirical evidence for the first time that artificially created landmarks in graphical command selection interfaces can improve the performance of spatial location learning and recall. The work introduced two types of artificial landmarks: abstract colour blocks providing clear visual anchors and semantically rich landmarking features of an image. With three progressively increased command set sizes, three studies showed that all artificial landmarks in GUIs provided spatial reference frames and supported spatial learning. For small interfaces, however, additional landmarks on top of grid lines offered little to no benefit. These studies also demonstrated that the importance of artificial landmarks gradually increased with the size of an interface.

Other contributions of this work include an empirical comparison of two different types of landmarking strategies. Although the image offered richer landmarks than the abstract colour blocks, the colour blocks appeared to provide better landmark benefits in spatial learning and retrieval of commands, particularly when the command set size increased. These new findings can assist designers when they build future graphical interfaces leveraging spatial memory.

4.4.2 Relevance in Context

In the context of my dissertation, Manuscript B reveals that the use of artificial landmarks in GUIs can support the learning and retrieval of command locations in command selection interfaces. It provided a confirmation that in the absence of adequate landmarks, digitally created *artificial landmarks* can provide support for spatial memory development of commands in GUIs. Since this study only considered command selection interfaces, a natural next step is to explore whether artificial landmarks are useful in other kinds of graphical interfaces. Besides command selection interfaces, widgets in graphical interfaces enable spatial locations revisitation in GUIs. For example, linear document viewers such as PDF viewers have 1D controllers called scrollbars allowing users to visit and revisit episodes in a document. Therefore, I carried out a study to investigate the effect of artificial landmarks in linear controllers (see Manuscript C in Chapter 5).

CHAPTER 5

EFFECT OF ARTIFICIAL LANDMARKS IN LINEAR DOCUMENT REVISITATION

Citation: Md. Sami Uddin, Carl Gutwin, and Alix Goguey. 2017. Using Artificial Landmarks to Improve Revisitation Performance and Spatial Learning in Linear Control Widgets. In *Proceedings of the 5th Symposium on Spatial User Interaction*, ACM, New York, NY, USA, 48–57.

Contributions and achievement: In coordination with Dr. Alix Goguey and my supervisor, Dr. Carl Gutwin, I designed and implemented the prototype interfaces for revisiting linear documents. In addition to directing this research, I was responsible for designing the study, data analyses, and reporting the study findings. Besides, with preliminary findings of this work, I designed and presented a poster[7] at the Graphics Interface - GI 2017 conference, which received the ***Best Poster Award***!

Besides command selection interfaces, learning and retrieving locations in GUIs is common in linear control widgets such as sliders and scrollbars that are widely used in navigating linear documents such as videos and text documents (i.e., PDFs). These linear controllers also involve spatial interactions because each of the locations on a controller represents a specific episode (e.g., a scene or a page) of a document. Therefore, Manuscript C aimed to investigate the effects of artificial landmarks on learning and recalling locations in individual widgets. I carried out a study with two types of linear controllers – a horizontal slider and a vertical scrollbar – where I compared artificial landmark augmented controllers with respective unadorned controllers.

[7] The poster that I designed with preliminary findings of a pilot study related to Manuscript B:
Md. Sami Uddin, Carl Gutwin, and Alix Goguey. 2017. "Artificial Landmarks Augmented Media Players for Video Revisitation" in *Graphics Interface (GI 17)*, Edmonton, Alberta.

5.1 PROBLEM AND MOTIVATION

Similar to command selection interfaces [50,76,87,88], GUIs for interacting with linear documents involve visiting and re-visiting episodes from a document. Linear control widgets (i.e., scrollbars or sliders) of these linear document interaction interfaces map episodes from a document (i.e., a video or a text) to locations on the widgets. Since these 1D controllers do not provide any visual guidance or landmarks for developing spatial memory of a document, carrying out efficient revisitation to specific locations within the document can be difficult and time-consuming [7].

In order to improve within-document navigation and revisitation facilities, researchers have explored different visual augmentations to the 1D controllers. For example, Hill et al. [100] projected the number of edits made to a text file on the scrollbars. Other researchers have augmented the 1D control widgets with different visual elements, such as alphabetic indices [3], information murals [152], thumbnails [67], and visit marks representing recently and frequently visited locations in a document [7]. Although these augmentations can improve document exploration and selection performance [7,198], they do not prove whether these techniques help users develop a spatial understanding of linear documents and support quick revisitation. Manuscript B in Chapter 4 indicated the value of artificial landmarks in command selection tasks at the level of an entire interface. Therefore, I saw an opportunity to investigate the effects of artificial landmarks in different contexts (at the level of an individual widget in linear document revisitation interfaces) and explore an obvious question – whether artificial landmarks support spatial learning in these 1D widgets.

The problem addressed in Manuscript C is that linear document interaction interfaces, particularly their 1D control widgets, do not provide enough landmarks to support efficient learning and remembering the spatial locations of episodes from a document.

5.2 SOLUTION AND STEPS TO SOLUTION

To address the lack of support for spatial memory development and efficient revisitation in linear document interaction interfaces, I investigated the use of artificial landmarks in linear widgets to figure out how they affect spatial learning and revisitation in different linear widgets. As most one-dimensional controllers are either horizontal or vertical, I tested both a horizontal slider and a vertical scrollbar. Additionally, I used two different kinds of content in my investigation: a video and a PDF document. Although both were linear documents, I saw an opportunity to examine the effect of artificial landmarks on learning and revisitation in two visually different contents: texts and videos.

In Manuscript B, I introduced two novel landmarking strategies for GUIs, but changes in the interfaces' layouts in Manuscript C (i.e., from command selection to linear document interaction) made those landmarks inappropriate for the linear widgets. Therefore, to test the effects of artificial landmarks on spatial location learning and recall in linear control widgets, in Manuscript C, I used two new landmarks: abstract icons and thumbnails of the content.

In the abstract icon strategy, I used random and monochrome abstract icons not related to the contents of the document to augment the one-dimensional controllers. The idea was that the spatial locations of those icons would represent different episodes of a document. Remembering those icon landmarks might enable users to perform quick revisits to previously visited document locations. As a second landmarking strategy, I used actual images extracted from the documents used in the study as thumbnails. Because of the space limitation in the controller, I manually selected a few thumbnails from the document, which could be easily differentiated from the rest.

5.2.1 Research Questions

This manuscript investigated the effects of two new artificial landmark augmentations in two types of linear controllers (horizontal sliders and vertical scrollbars) for revisitation tasks. I ran a study to compare these new landmark techniques with two non-landmark versions, each for one of the two controllers. The study focused on answering the following questions:

- Do artificial landmark augmentations support spatial learning in linear widgets?
- How do artificial landmarks affect linear document revisitation performance?
- What type of artificial landmarks support better spatial learning?
- Do the content of the document and a target's proximity to a landmark affect revisitation performance?

5.3 MANUSCRIPT C

Linear interface controllers such as sliders and scrollbars are primary tools for navigating through linear content such as videos or text documents. Linear control widgets provide an abstract representation of the entire document in the body of the widget, in that they map each document location to a different position of the slider knob or scroll thumb. In most cases, however, these linear mappings are visually undifferentiated – all locations in the widget look the same – and so it can be difficult to build up spatial knowledge of the document, and difficult to navigate back to locations that the user has already visited. In this paper, we examine a technique that can address this problem: artificial landmarks that are added to a linear control widget in order to improve spatial understanding and revisitation. We carried out a study with two types of content (a video and a PDF document) to test the effects of adding artificial landmarks. We compared standard widgets (with no landmarks) to two augmented designs: one that placed arbitrary abstract icons in the body of the widget, and one that added thumbnails extracted from the document. We found that for both kinds of content, adding artificial landmarks significantly improved revisitation performance and user preference, with the thumbnail landmarks fastest and most accurate in both cases. Our study demonstrates that augmenting linear control widgets with artificial landmarks can provide substantial benefits for document navigation.

5.3.1 Introduction

Linear documents such as text, webpages, audio, video, or slideshows typically use linear interface widgets for navigation (e.g., scrollbars or sliders). These widgets provide an abstract spatial representation of the entire document (although not a visual representation), in that one dimension of the controller is absolutely mapped to document length (e.g., Y-position of a PDF, or timestamp in a video). Aside from this spatial mapping, most linear controllers do not provide any visual marks that represent document content. Often linear controllers (e.g., YouTube or Adobe Acrobat Reader) provide interactive thumbnails showing only a small portion of the document (but not the whole document), allowing users to access the content close to the focused region. As a result, using the controller to develop a spatial understanding of the document, and to remember and revisit specific document locations, can be difficult [7].

Several researchers have proposed visual augmentations to one-dimensional controllers to address a variety of navigation problems. For example, some techniques show a document's interaction history, such as Hill et al.'s edit-wear scrollbars that showed the number of edits in a text file [100]; other visualizations show notifications such as the location of syntax errors in a code editor. A few of these projects have used augmentations that can support the development of spatial memory – such as the AlphaSlider's alphabetic index [3], the document maps of the Mural Bar [152], Code Thumbnails [67], or the video thumbnail grids of the Swifter video scrubber [151]. However, although studies have shown that these augmentations can improve search and selection tasks [7,198], there is little evidence about how well they support spatial understanding and revisitation. The exception is a technique called the Footprints scrollbar [7], which visualizes recently-visited and frequently visited document locations. A study showed that the Footprints scrollbar aided navigation back to previous locations, and that showing ten marks in the scrollbar could account for a substantial proportion of revisits [7].

A problem with techniques based on visit histories, however, is that there are many situations where a user may want to revisit a location not shown in the widget. For example, if a user watches an entire video clip, all locations are visited equally during the initial playback, making it harder to go back to a particular scene. Similarly, many visited locations may not appear in a Footprints-style augmentation: for example, some locations have not been seen often enough (or for long enough at each visit) to appear in the visualization, and some locations that go unvisited can disappear from list of recent items.

These problems arise because visit-history techniques depend on the system to remember and visualize important document locations. A different approach is to rely on the users to remember these important locations – which is possible if they are given the resources to exploit their spatial memory abilities. Human spatial location memory is highly effective, and can be both expansive and accurate if the environment is rich and spatially stable (e.g., [7,183,218]). Therefore, it is possible that revisitation with linear widgets can be substantially improved simply by adding a rich set of spatially-stable landmarks to the controller – allowing users to build up spatial memory of important document locations.

In this paper, we report on a study of how artificial landmarks affected spatial learning and revisitation using two different linear widgets (a horizontal slider and a vertical scrollbar) and two different kinds of content (a video and a PDF document). The study asked participants to find and then revisit different locations in the documents. For each system, we compared a standard widget with no landmarks, a widget augmented with a set of arbitrary abstract icons; and a widget with thumbnails extracted from the content. Our results showed that both kinds of artificial landmark improved users' spatial learning and revisitation performance, with the thumbnail condition performing best in both systems.

Our work provides three contributions. First, we demonstrate two designs for augmenting linear control widgets with artificial landmarks. Second, we show that artificial landmarks can significantly improve revisitation performance, in two different contexts. Third, we provide new empirical evidence that adds to our understanding of spatial learning in user interfaces, and that helps to confirm users' ability to learn and navigate documents using spatial memory.

5.3.2 Related Work

5.3.2.1 Revisitation in User Interfaces

Prior work has shown that human behaviour with interactive systems is highly repetitive [117,236] – although systems often present several ways to complete a task, users mostly choose mechanisms that they are familiar with. Revisitation is one kind of repetitive behaviour in which users return to the same locations over and over, a phenomenon that has been most clearly established in the use of the web [2,53,206]. Revisitation patterns have also been observed in menu selections [50,76,87,88], document readers [7], and video players [32], and researchers have looked at several aspects of how revisitation works and how it can be supported with interface augmentations.

Prior work on revisitation support can be divided into two groups. First, there are manual techniques which rely on explicit user actions. The most common tools falling into this category are bookmarks, which are widely available in traditional web browsers, document viewers and office applications [7]. Bookmarks allow users to manually set flags to simplify future revisits of particular content. This idea was used in the Bookmark Scrollbar [131], which places bookmarks in a classic scrollbar. There are, however, some limitations associated with bookmarks' use in interfaces [1]: first, the user has to recognize that a particular location will be revisited (which may not be apparent at the time); second, users must manually place bookmarks, which can discourage users from doing so; and third, people sometimes persist with suboptimal strategies [197] even if a more efficient long-term solution is available [1] (e.g., browsing again through a document vs. using bookmarks).

The second type of revisitation support involves automatic techniques that do not require manual interventions from the user. Interface controls such as ‘Forward/Back’ buttons and ‘Recent Items’ menus monitor interactions in the background [7] and update the interface accordingly. However, as observed by Alexander et al. [7], people often misunderstand the use of the ‘Recent Documents’ menu and the ‘Back’ buttons in typical web browsers [7,53], causing problems for efficient revisitation. Other automatic techniques are more explicitly focused on revisitation – such as Hill et al.’s [100] ‘read wear’ that shows histograms of a user’s reading history, Alexander et al.’s [7] Footprint Scrollbar that adds transient history marks in the scrollbar, and Skopik and Gutwin’s ‘visit wear’ augmentations that added explicit visit marks to a fisheye visualization system [198].

Much of the previous work on revisitation has involved text documents, but some research has been carried out on navigation in video players. Most media players support approximate revisitation with ‘Forward’ and’ Rewind’ buttons, but precise revisitation remains a problem. Matejka et al. [150,151] tried to improve video scrubbing by showing thumbnails of the video frames over the slider when the cursor is on the slider knob (a technique also seen in some recent video players such as YouTube – showing only partial range of scenes as thumbnails).

5.3.2.2 Interface Augmentation

Interface augmentation is a widely-accepted method to improve computer systems expressivity and interactivity. In GUIs, it is common to augment interfaces with colors, symbols, images or icons, and some researchers have used these techniques to improve document revisitation. As discussed above, ‘edit wear’ and ‘read wear’ techniques [100] show interaction histories, and many code editors augment the scrollbar with annotations (e.g., about syntax errors [7]). Other document navigation systems have used colored marks [7,41] or visualizations of content [3,67,152] to augment scrollbars and improve search (e.g., Code Thumbnails, Mural Bar, or AlphaSlider).

Media players are another type of interface where augmentation is commonly used, although previous work does not appear to explicitly focus on revisitation. To support exploration within a video, researchers augmented the timeline slider of the media player, showing visual highlights to represent personal [5] or crowd [124,231] navigation history. Chen et al.'s [47] Emo Player annotates the media player's slider with different colours based on the characters' emotional states. Other techniques show interactive thumbnails from the video as a storyboard on the screen [35,110], giving users an overview of the entire video with a large grid of thumbnails which can assist exploration. Instead of extracting entire frames from the video, Schoeffmann et al.'s video explorer [190] augments the slider with the dominant colours of each frame to help users navigate and explore video content.

5.3.2.3 Use of Artificial Landmarks in Interfaces

In GUI-based systems, landmarks such as the corners of the screen can provide a strong external reference frame that helps users build spatial memory [215,217]. Several techniques have explicitly made use of the bezel and corners of small devices (e.g., tablets or smartwatches) as landmarks to organize menus and toolbars [95,132,191]. In areas where there are few natural landmarks (such as the middle of a large screen), artificially placed objects (e.g., coloured blocks) or even the user's own hands [217,220] can act as landmarks and help users to navigate through the interface and recall command locations [7,218].

This previous work indicates the usefulness of landmarks (even artificial ones), for improving document navigation and command selection. In the next section, we explore the use of artificial landmarks explicitly in linear document navigation.

5.3.3 Artificial Landmarks for Linear Document Controls

To test the performance of artificial landmarks on revisitation in linear documents, and to see how different kinds of augmentation affect performance, we designed new variants of standard slider and scrollbar widgets that are augmented with landmarks. Scrollbars are used to navigate through text, pictures, or any other content in a predetermined direction (vertical or horizontal), when the display can only show a fraction of the content at once. A slider is used to set or pick a value by moving an indicator along a defined segment, usually in the horizontal direction. These widgets can be manipulated easily with a mouse in a standard desktop interface or with a finger in a touch interface.

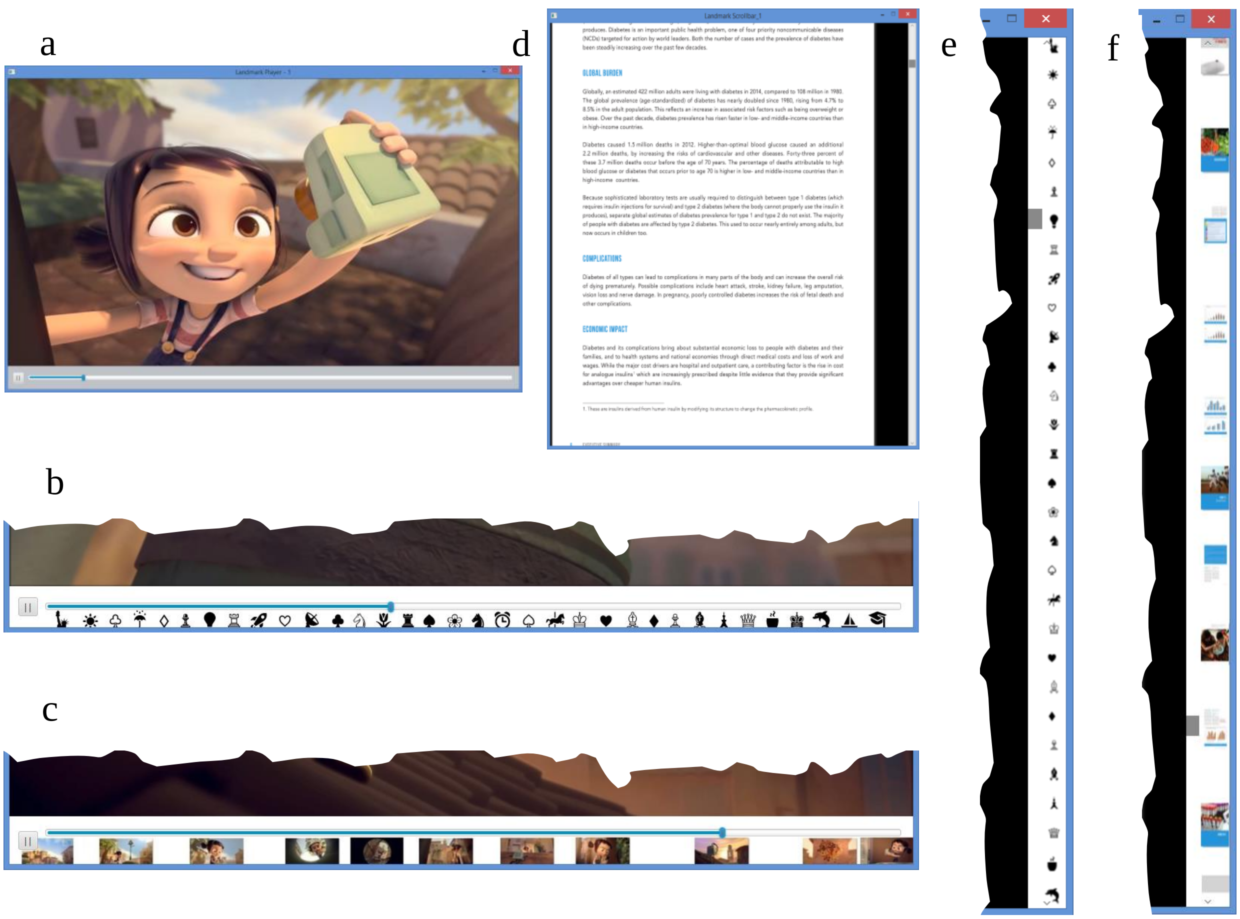

Figure 5.1: Study interfaces. Media player (a, b, c), PDF viewer (d, e, f). A, d: standard - with no landmarks; b, e: icon – augmented with abstract icons; c, f: thumbnail – augmented with extracted content as thumbnails. Sources [247,248].

We developed two versions of each augmented widget, in addition to the standard version. Figure 5.1 shows our three media player interfaces: all are 1200×726px in size and all have a 1D linear control (60px tall) at the bottom of the window. All versions support the same navigation features (play/pause button and timeline slider); the only difference between the interfaces is in the augmentation of the control widget. The three versions of the slider were: *standard* (the ordinary slider with no augmentation), the *icons* version with abstract icons as landmarks that had no relation with the contents, and the *thumbnails* version with thumbnails extracted from the video (see Figure 5.1: a, b, c). When the user presses the play button, the video starts, and the user can click on the slider to go to any desired location of the video. We removed 'scrubbing' functionality from all three sliders to allow us to better record the users' location choices (i.e., all navigation actions were through clicking on the slider or using the play button).

Our three versions of the PDF viewer (see Figure 5.1: d, e, f) also have the same layout and size (900×890px), with a 1D control at the right of the window. Our three versions were similar to those described above, with *standard* (40px wide), *icons* (50px wide), and *thumbnails* (60px) widgets. All versions of the PDF viewer allowed users to view and navigate through the pages of a document only by clicking on the scrollbar. Clicking anywhere on the scrollbar immediately takes the user to the corresponding page; as with the video player, interactive dragging was turned off to get a more accurate measure of users' spatial location memory.

Standard

The standard versions do not provide any extra landmark other than those naturally embedded in a regular slider or scrollbar. The relative position of the thumb in the controller can be used to infer the location within the entire document. See Figure 5.1: a and d.

Icons

As the length of the content increases, it is more likely that the relative location cue provided by the scrollbar thumb will become less effective in aiding spatial memorization and retrieval. The icons interface therefore augments the widget with monochrome abstract icons (arbitrary; unrelated to the contents) that are distinct from their surroundings [222], and that provide clear spatial reference points [218]. We placed 34 icons (each 26px in size) horizontally in the media player, and 30 icons (22px each) vertically in the document viewer. See Figure 5.1: b and e.

Thumbnails

We augmented the control with actual images extracted from the content. Because we have limited space in the control, we selected thumbnails based on important scene transitions (in the video) or visually distinct pages (in the PDF document). We also kept the inter-thumbnail distance (e.g., time interval in video and number of pages in PDF) approximately uniform. We used 11 thumbnails (each 70x40px) horizontally in the media player and 11 thumbnails (each 32x42px) vertically in the document viewer (see Figure 5.1: c and f. There are also methods to automatically extract key elements from documents [168] or video [167], which could be used in future systems.

5.3.4 Study: Effects of Artificial Landmarks on Revisitation

We ran a study to examine revisitation performance with our three versions of the linear control widgets. We designed the study to answer two main questions: first, do artificial landmarks improve revisitation compared to the standard widgets; and second, do extracted thumbnails perform better than abstract icons.

5.3.4.1 Study Methods and Design

To ensure that the three widget designs were fairly compared for each content type, we used the same video and document for all three widgets. This meant that we could not have participants complete tasks with all three versions of the interface, as they would have built up experience from one interface to the next. Therefore, we chose a mixed within-participants / between-participants design for the study. Each participant used both the video player and the PDF viewer, and used a different widget design for each system. This meant that there were three groups:

- G1 (10 people): *standard* and *icons*
- G2 (10 people): *icons* and *thumbnails*
- G3 (10 people): *thumbnails* and *standard*

All groups were counterbalanced so that five people used each interface in each system (e.g., in G1, 5 people used *standard* for the video player and 5 used *icons)*. We included group as a factor to check for grouping effects; because there were none (as described below), we carried out our comparisons using all ten people in each system+interface combination.

5.3.4.1.1 <u>Media Player: Tasks and Stimuli</u>

Participants started by watching a video [247] twice using one of the three custom media player interfaces (standard, icons, or thumbnails) that ran on the right screen of a dual-monitor (21-inch) environment. Each participant then went through a series of trials where they were asked to navigate to a specific target frame by clicking on the slider. Each trial began by displaying the stimulus frame on the left screen. The participant then used a mouse to locate the target frame. To make a correct revisitation, participants had to click within 10px (which corresponds to 30 frames) of the target's actual location on the slider.

The video (length 2:12) shown in the study was the same across all three media players as each participant used only one version of the media players. Eight frames from the video were manually selected and used as stimuli. They remained the same across all conditions. In the thumbnail condition, the locations of three out of the eight selected target stimuli were on the visible landmarks, and the rest were between or near to the landmarks. We also made sure that the selected targets were spaced regularly throughout the video.

5.3.4.1.2 Document Viewer: Tasks and Stimuli

The PDF tasks were similar to those described above. Participants had two minutes to become familiar with a 42-page PDF report [248] using one of the three interfaces (standard, icons, or thumbnails). Each participant then completed a series of trials where they were asked to navigate to a specific target page by clicking on the scrollbar. Each trial began by displaying the stimulus page. The participant then used a mouse to locate the target page. To make a correct revisitation, participants had to click within 12px (which corresponds to half of a target page's height) of the target's actual location on the scrollbar. We used eight pages (spaced approximately regularly through the document) as stimuli. In the thumbnail condition, three target pages were located on the visual representation of a landmark.

5.3.4.1.3 Procedure and Study Design

We explored two kinds of content (*media player* and *PDF viewer*) and analysed results separately for each system. For each system, we analysed the effects of interface condition (*standard*, *icons* and *thumbnails*) on revisitation time and errors. Each participant used both systems and saw two different interfaces (but only one version of each interface for each system, as described above). The order of applications and conditions was counterbalanced.

Participants were instructed to complete the trials as fast and as accurately as possible. When using the video system, participants had 15 practice revisitations using a different video than the one used in the main experiment. Participants then completed 5 blocks of trials (each consisting of the same 8 target stimuli, presented in random order). After completed all blocks, participants completed a NASA-TLX subjective workload questionnaire [96]. When using the PDF application, participants had 15 practice revisitations using a different document, and then completed 5 blocks of trials (each consisting of the same 8 stimuli, presented in random order). The PDF task was also followed by another NASA-TLX questionnaire. Participants started with one of the two applications and then proceeded to the other.

For each trial, a selection on the slider or scrollbar displayed the corresponding video frame or PDF page. Participants could adjust their selection up to ten times for each trial. Our software recorded trial completion time, errors, and data describing every selection. At the end of the study, participants provided subjective responses through a questionnaire.

5.3.4.1.4 Participants and Apparatus

Thirty participants (8 female), ages 19-30 (mean 24.7), were recruited from a local university. The study took 30 minutes on average. Each participant was compensated with a $5 honorarium.

The experiment was conducted on a desktop computer running Windows 8.1, with two 21-inch 1920x1080 resolution monitors placed alongside. Software was written in JavaFX. Input was received through an optical mouse. All study interfaces ran centered in the right screen (with a white desktop background).

5.3.5 Study Results

We report the effect size for significant between subject RM-ANOVA results as partial eta-squared: η^2 (considering .01 small, .06 medium, and >.14 large [30]), and Bonferroni correction was performed for post-hoc t-tests.

Before starting our analyses, we first checked for the effect of our grouping variable (see Section 5.3.4.1). ANOVA showed no effect of *group* (for the video player, F=0.77, p=.47; for the PDF viewer, F=2.62, p=.09), so we conducted further analyses using all participants for each system+interface combination.

5.3.5.1 Media Player: Results

For the media player tasks, 28 out of 1200 trials were discarded from analyses (either because completion time was more than two s.d. away from the respective mean of each block, or the trial could not be completed within 10 attempts).

5.3.5.1.1 Media Player: Trial Time

Mean completion times across blocks for the three conditions are summarized in Figure 5.2. RM-ANOVA showed a significant main effect of *condition* ($F_{2,27}$=4.45, p=.02, η^2=0.2): 7746ms (s.d. 5804ms) for *thumbnail*, 10306ms (s.d. 5809ms) for *icon*, and 11588ms (s.d. 6703ms) for *standard.* Completion times decreased across *block* ($F_{4,108}$=54.86, p<.001, η^2=0.34), and as anticipated, the skill development follows a power-law function [159]. There was no significant *condition* × *block* interaction ($F_{8,108}$=0.94, p=.48).

Post-hoc pairwise t-tests (Bonferroni-corrected) showed that both artificial-landmark conditions were faster than *standard* (all $p < 0.01$) but showed no difference between *thumbnails* and *icons.*

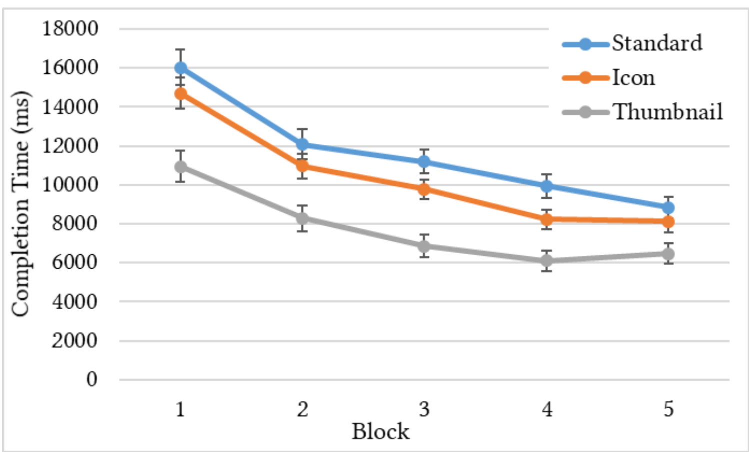

Figure 5.2: Mean completion time by block and interface.

5.3.5.1.2 Media Player: Target Proximity to Landmark

We also analyzed the trial time for *thumbnail* condition based on the proximity of the target to a landmark (3 of 8 targets were located on the thumbnails). As shown in Figure 5.3, the on-thumbnail targets were faster than near-thumbnail targets, with mean completion time of 6761ms (s.d.

5366ms) and 8315ms (s.d. 5979ms). However, ANOVA did not show any significant differences between these two categories ($F_{1,18}$=0.51, p=.48).

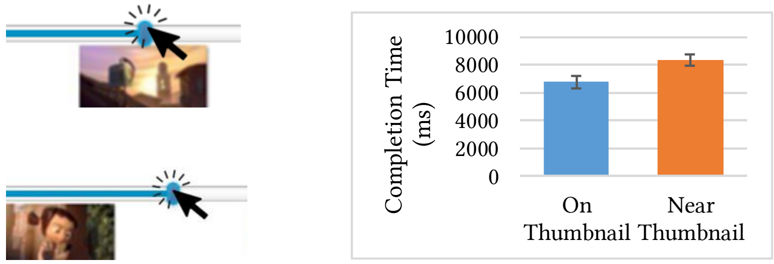

Figure 5.3: On-thumbnail vs near-thumbnail selections. Left-top: on thumbnail, Left-bottom: near thumbnail. Right: mean completion time.

5.3.5.1.3 <u>Media Player: Error Rate – Overall</u>

The number of errors per trial (i.e., the number of attempts to complete a trial) is summarized in Figure 5.4. There was a significant main effect of *condition* ($F_{2,27}$=4.27, p=.02, η^2=0.17), but no *condition* × *block* interaction ($F_{8,108}$=1.1, p=.37). Post-hoc t-tests (all p<.001) showed *thumbnail* had the lowest error rate at 2.84 errors/trial (s.d. 3.43), compared to higher error rates of 4.21 (3.37) for icon and 4.30 (3.37) for *standard*. As expected, error rates were higher at the beginning, and decreased significantly with *block* ($F_{4,108}$=8.94, p<.001, η^2=0.1).

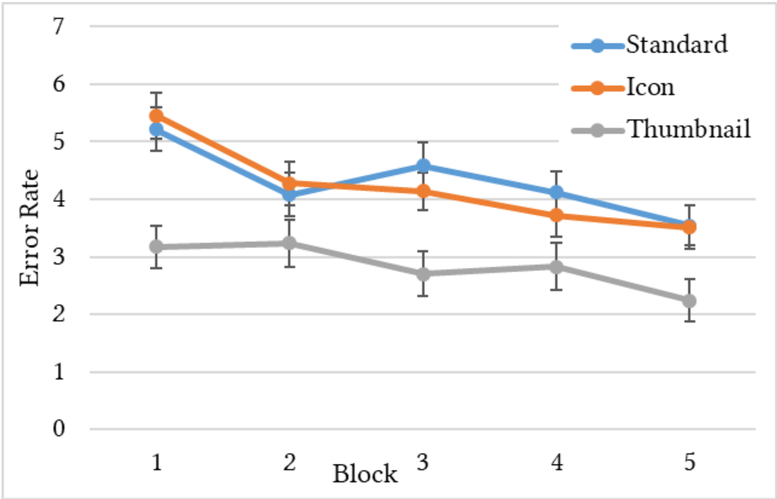

Figure 5.4: Error rate by block and interface condition.

5.3.5.1.4 Media Player: Analyses by Target

We also analyzed errors for each target (stimuli were the same in all three conditions). ANOVA showed a significant difference of *target* ($F_{7,216}$=19.04, p<.001, η^2=0.38) and *condition* ($F_{2,216}$=11.35, p<.001, η^2=0.1), but no *target* × *condition* interaction ($F_{14,216}$<1). Figure 5.5 shows the errors/trial for each target and their location in the video. ANOVA also showed significant completion time differences for *target* and *condition* (all F>7, p<.001, η^2>0.2). Post-hoc t-tests showed that the targets placed at the middle (especially targets 5 and 7 for *standard* and *icon* took more time (Figure 5.5; all p<.02) and was error prone (targets 4, 5 and 7; all p<.03), following the serial position effect [55]. However, *thumbnails* was significantly better than *standard* for the middle targets (2 and 4-7, all p<.01).

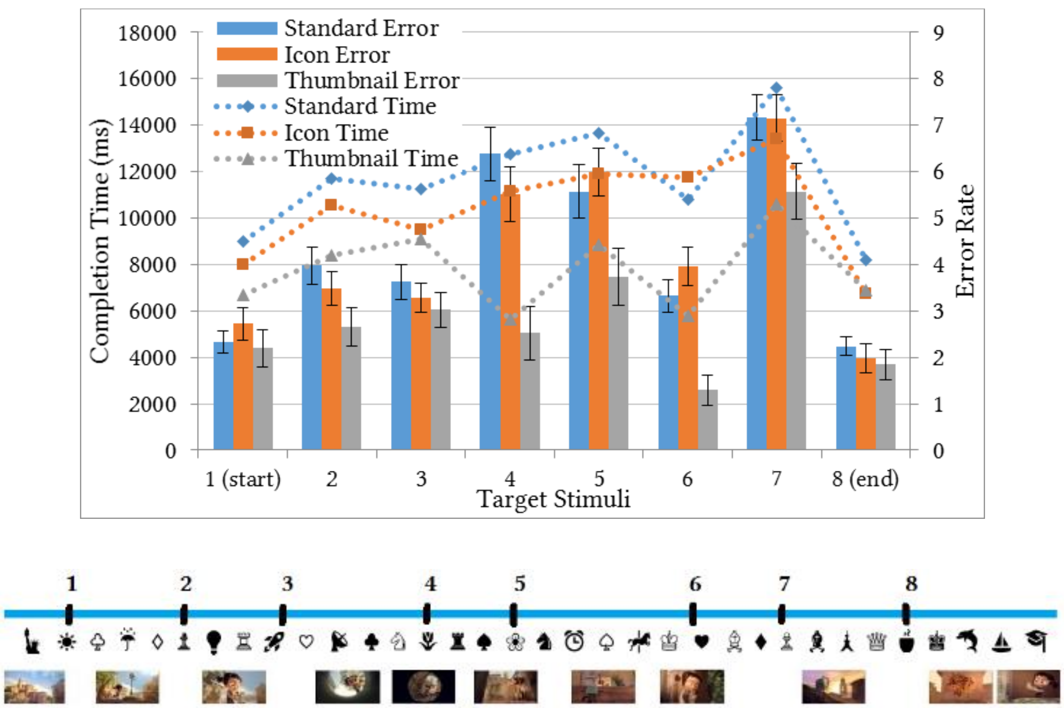

Figure 5.5: Analyses by target. Top: results, Bottom: locations of the target stimuli in the slider.

5.3.5.1.5 Media Player: Subjective Responses

Participant responses on the NASA-TLX worksheets showed significant differences for the three conditions (Kruskal-Wallis tests, Table 5.1: low score means better, except for Performance). Overall, both landmark conditions performed well. Post-hoc t-tests showed that *thumbnails* achieved significantly better scores for Mental, Physical, Performance, and Frustration scales (all $p<.02$).

Table 5.1: Mean (s.d.) effort scores (0-10 scale, low to high).

	Standard	Icon	Thumbnail	χ_r^2	*p*
Mental	8.1(1.2)	5.1(2.5)	5.2(2.44)	9.92	**.01**
Physical	7.0(1.76)	2.6(1.71)	3.9(1.85)	16.13	**.01**
Temporal	6.3(1.06)	4.7(1.83)	4.1(1.66)	8.22	**.02**
Performance	5.8(1.81)	6.7(2.41)	7.0(2.0)	1.8	.41
Effort	7.7(1.42)	5.5(2.46)	6.1(2.08)	4.97	.84
Frustration	7.0(1.56)	4.4(2.12)	4.0(2.26)	10.55	**.01**

5.3.5.2 PDF Viewer: Results

For the PDF viewer, 34 out of 1200 trials were discarded from analyses (completion time was more than 2 s.d. away from the mean, or the trial not be completed within 10 attempts).

5.3.5.2.1 PDF Viewer: Trial Time

Mean completion times across blocks for the three conditions are summarized in Figure 5.6. RM-ANOVA showed a significant main effect of *condition* ($F_{2,27}$=7.88, p=.002, η^2=0.3): (5236ms, s.d. 3733ms) for *thumbnail*, (6820ms, s.d. 5666ms) for *icon*, and (8685ms, s.d. 5618ms) for *standard*. As with the video player, there was a significant effect of *block* ($F_{4,108}$=53.76, p<.001, η^2=0.35), but no *condition* × *block* interaction ($F_{8,108}$=1.13, p=.35).

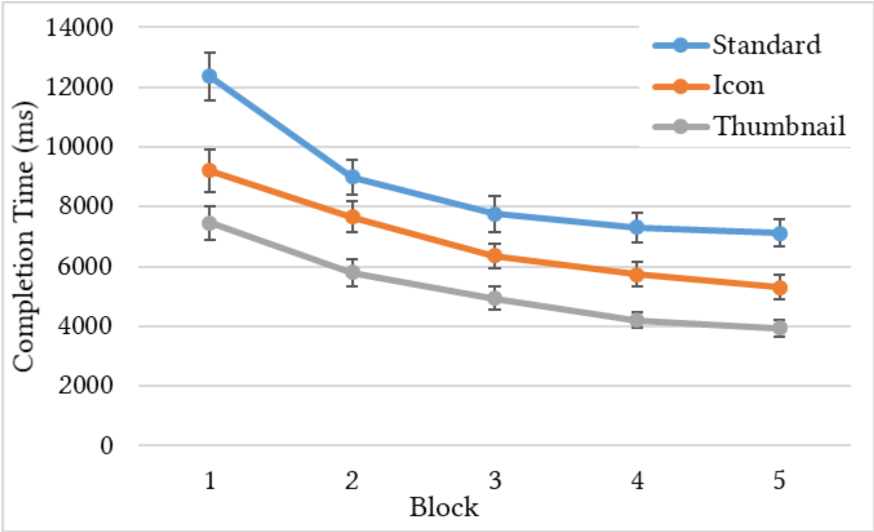

Figure 5.6: Mean completion time for the three PDF viewers.

5.3.5.2.2 PDF Viewer: Target Proximity to Landmark

Unlike the video, ANOVA showed a significant difference for the PDF viewer between on-thumbnail and near-thumbnail targets ($F_{1,18}$=9.12, p<.01, η^2=0.34), with 3859ms (s.d. 2698ms) for on-thumbnail and 6372ms (s.d. 4419ms) for near-thumbnail (Figure 5.7).

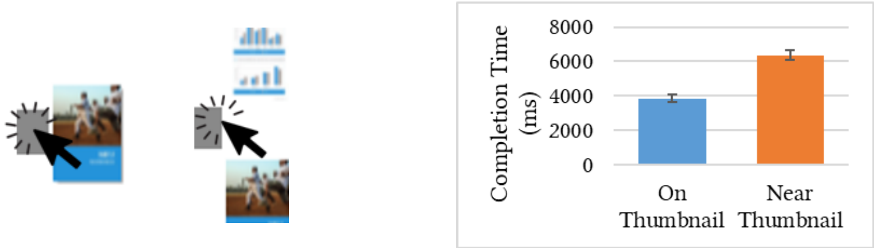

Figure 5.7: On-thumbnail (left) vs near-thumbnail (middle) selections. Right: mean completion time.

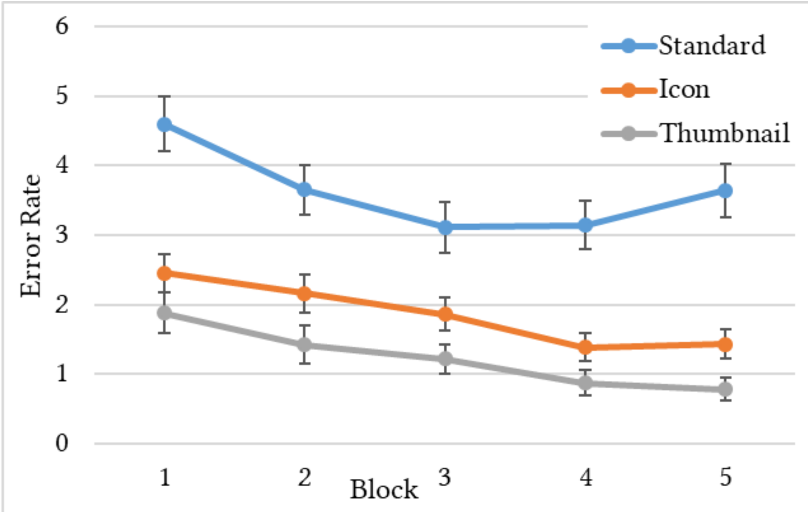

Figure 5.8: Error rates for document viewers.

5.3.5.2.3 PDF Viewer: Error Rate – Overall

There was a significant main effect of *condition* on errors ($F_{2,27}$=24.53, p<.001, η^2=0.51): *thumbnails* had the lowest error rate at 1.22 errors/trial (s.d. 2.03), compared to 3.62 (3.3) for *standard* and 1.85 (2.15) for *icons*. As shown in Figure 5.8, errors were high in the first blocks and generally decreased over time (except in the final blocks of the *standard* condition), leading to a significant effect of *block* ($F_{4,108}$=12.21, p<.001, η^2=0.16), but no *condition* × *block* interaction ($F_{8,108}$=0.94, p=.49). Post-hoc t-tests showed that all three interface conditions were significantly different (all $p < 0.01$).

5.3.5.2.4 PDF Viewer: Analyses by Target

Figure 5.9 shows the error rates by targets for all conditions in the PDF viewer. ANOVA showed a significant difference among *targets* ($F_{7,216}$=7.73, p<.001, η^2=0.2) and *conditions* ($F_{2,216}$=54.5, p<.001, η^2=0.34). There was also a significant interaction between *targets* × *conditions* ($F_{14,216}$=3.26, p<.001, η^2=0.17).

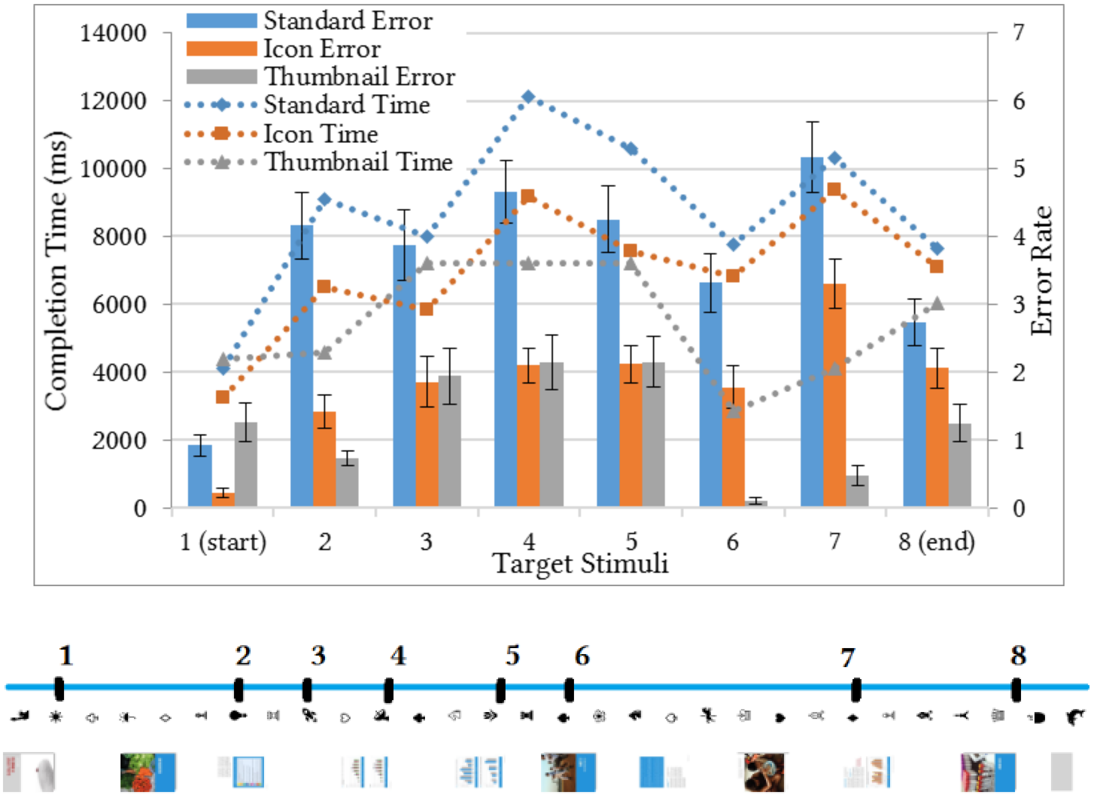

Figure 5.9: Analyses by target. Top: results, Bottom: locations of target stimuli on the scrollbar.

Targets from the middle areas were most error-prone for *standard* (post-hoc analysis for targets 2-7: >3.32 errors/trial, all p<.001); in contrast, landmarked versions performed better throughout, including the middle areas. ANOVA also showed significant completion time differences for *target* and *condition* (all F>11, p<.001, η^2>0.1). T-tests showed that the landmarked conditions outperformed *standard* at the middle (targets 2, 4 and 5 all with mean completion times <9188ms, p<.05; Figure 5.9).

5.3.5.2.5 Subjective Responses

Table 5.2 summarizes mean responses to the NASA-TLX worksheets (low score means better, except for Performance). Kruskal-Wallis tests and post-hoc t-tests (all p<.03) showed significant effects for Mental, Temporal workload and Frustration (*thumbnail* lowest, *standard* highest), and for Performance (*thumbnail* highest, then *icon* and *standard* lowest).

Table 5.2: Mean (s.d.) effort scores (0-10 scale, low to high)

	Standard	Icon	Thumbnail	χ_r^2	*p*
Mental	6.8(1.48)	5.7(1.89)	4.4(1.65)	7.9	**.02**
Physical	5.4(2.55)	4.4(2.55)	3.6(1.51)	2.68	.26
Temporal	5.5(0.85)	4.3(1.06)	3.1(1.2)	14.77	**.01**
Performance	6.5(1.58)	7.0(1.63)	8.3(0.95)	6.95	**.03**
Effort	6.1(2.51)	5.7(1.34)	4.8(1.32)	2.5	.29
Frustration	5.5(2.42)	4.4(1.84)	1.8(1.48)	12.49	**.01**

5.3.5.3 Participant Preferences and Comments

After completing both studies, participants provided their preferences between the two conditions they used. Table 5.3 shows that participants favoured both artificial-landmark conditions, with 80% of participants preferring them across five measures. The *thumbnails* interfaces were preferred by 70% of participants between the two landmarked conditions.

Table 5.3: Count of participant preferences.

	Group 1		Group 2		Group 3	
Abstract	Standard	Icon	Icon	Thumb	Thumb	Standard
Speed	2	8	2	8	8	2
Accuracy	2	8	3	7	7	3
Memorization	2	8	2	8	10	0
Comfort	3	7	2	8	8	2
Overall	2	8	3	7	8	1

Participant comments echoed our other findings. Participants made several comments on how the landmarks (especially thumbnails) helped them to develop spatial memory of the contents. One participant stated, *"Snapshot of a specific action [thumbnail] helped me remember the story sequence [of the video]."* Another mentioned *"I remembered the sequence [of the video] after seeing the closest thumbnail."* For the thumbnails in the PDF viewer, one person mentioned *"Thumbnails [of the pages] made it easy to find exact pages and guess the nearby page locations."*

Other comments indicated that the icons also helped participants to learn command locations: one mentioned *"It was easy to remember which page was beside which icon [in the PDF viewer]."* Some participants revisited locations by correlating icons with the content: as one said, *"I could correlate some of the icons with video contents."* A few participants, however, found the icons more difficult: *"The randomness of the icons was a little tough for me to remember [in PDF viewer];"* another said, "*I associated a few key locations [of the PDF document] to icons [but] it was difficult to keep track of the icons*."

5.3.6 Discussion

Previous research has identified that revisitation is common in computer interfaces [7,32], and that artificial landmarks can be useful in helping users remember item locations for future visits [215,218]. Our study results suggest that spatially-stable artificial landmarks can help people to remember locations and can improve revisitation performance. Our study provides two main results:

Both artificial-landmark interfaces were faster than the standard widget for both applications.

Of the two types of artificial landmarks, the *thumbnails* condition was fastest and most accurate, and was strongly preferred by participants.

5.3.6.1 Explanation and Interpretation of Results

5.3.6.1.1 <u>Landmarked Interfaces can Improve Spatial Revisitation</u>

During the study, mean completion time decreased significantly across blocks for all conditions in both applications; and as anticipated, they followed the power-law of learning curve [159]. The number of errors also decreased across blocks. These can be indications that users transitioned quickly from slow visual search revisitation to more rapid spatial-memory-based revisitation. This transition occurs in all conditions for both applications, confirming that users developed spatial memories, regardless of the interface. However, Figures 5.2, 5.4, 5.6 and 5.8 reveal that participants were successful in forming and exploiting these memories more rapidly and accurately when interfaces were augmented with artificial landmarks (*icons* and *thumbnails*) for both applications.

Figures 5.2 and 5.6 indicate that although the *standard* interface for both applications allowed participants to build a certain degree of spatial memory of the contents, participants were less able to rely on their spatial memory to perform rapid actions [183,218]. As shown in Figures 5.4 and 5.8, participants in the *standard* condition made substantially more mistakes in revisitation, especially in the later blocks of the PDF viewer, likely because the standard widgets provided no clear reference frame. As a result, *standard* interfaces with no landmarks were slowest and most error prone.

5.3.6.1.2 Why did Thumbnails Outperform Icons?

Our results show that in both applications, several measures favoured the *thumbnails* interface over *icons* (trial time, error rates, workload measures, and preference). We see two key factors responsible for this advantage. First, the thumbnail landmarks in the *thumbnails* interfaces were actual representations of the contents as images in miniature scale. The thumbnails showed a clear mapping between the control widget and the content, compared to *icons* more abstract icons that had no connection with the actual content. Participants had to manually form the mappings between icons and contents for later visits. Though participants were successful in learning the mappings and revisited target stimuli, additional time was required to learn the connections, which may explain the slower performance of *icons* compared to *thumbnails*. The clearer indication of document content may have helped participants to exploit their spatial memory and perform rapid revisitations with better accuracy. Performance of the *icons* condition could potentially be improved by choosing icons that are more meaningful to the content (although this may not always be possible).

Second, the number of landmarks available in the interfaces may have caused the performance difference between the two interfaces. *Icons* interfaces had more landmarks (more than 30) compared to *thumbnails* (11). Overloading the 1D controller of an interface with more items means users require to learn and remember more mappings. Learning and remembering only 11 items may have been easier – determining the ideal number of landmarks is an interesting area for future study. Additionally, the available contents of the document, especially in the PDF document where colourful charts and images were present in more than 60% pages, could have helped users to form spatial memory. Exploring the influence of document-contents (e.g., pages with images or only texts) on spatial memory formation is a compelling area of future work.

5.3.6.1.3 Landmarks can Help Overcome the Serial Position Effect

The serial position effect [55] suggests that people best recall the first (primacy effect) and the last (recency effect) items in a series compared to the middle items [66,158]. During our study, we showed the entire content (video and document) to the participants first, then asked them to revisit or recall items shown as target stimuli. Analyses by targets (see section 5.3.5.1.4 Figure 5.5 and section 5.3.5.2.4 Figure 5.9) show that standard interfaces for both applications followed the serial position effect, meaning that the targets used from the beginning and the ending of the documents were recalled more accurately than the middle targets. However, we see that landmark-augmented interfaces appeared to overcome this effect: participants were able to revisit targets in the middle with better accuracy in both applications, especially with *thumbnails*. We believe this was possible because of the presence of the artificial landmarks in *icon* and *thumbnail* interfaces. Landmarks provided spatial anchors for the participants that helped to remember middle area's contents and allowed them to revisit those locations accurately.

5.3.6.2 Implications for Design

Our findings provide additional evidence that artificial landmarks can assist in developing users' ability to build spatial memories for rapid revisitation [7,218]. Many interactive interface components often limit users from utilizing the full potential of their spatial memory because they lack landmarks. For example, the blank trough of a slider or scrollbar may lead to inefficient revisitations. Past research has attempted to temporarily augment scrollbars with interaction histories [7,100] or usage information of a video in media player's slider [124]. Our results suggest that static embellishments of the interface with artificial landmarks can substantially improve users' revisitation experiences.

There are, however, a few challenges involved in augmenting interfaces with landmarks. First, there is a chance of interference from artificial landmarks in the interface and from the primary content of the interface. While the user is focusing on the main content, the static landmarks on the controller might catch unwanted attention from the user. Similarly, focus on the main content might push the landmarks out of the user's focus. Careful design of the landmarks and proper placements can overcome this issue. For instance, landmarks can be highly transparent to reduce unwanted focus; this way, users only switch their attention on the landmarks when they need to remember content. Further studies are needed to explore real-world issues in the use of artificial landmarks.

A second challenge is the length of the content (e.g., video duration or document length). For a relatively large document, small movements of the slider/scroller result in large navigation actions. A possible solution to this problem is adding additional controllers – each controller could be responsible for specific range (e.g., 30-minute video or 50-page document) of the full content, and still provide landmarks to support revisitation.

Third, the success of the *thumbnails* interface may be related to our ability to choose appropriate and meaningful thumbnails from the document content. Although previous work in automatic extraction of summary information from a video has proven to be successful (e.g., [167,168]), a clear direction for future work is to replicate our study using an algorithm for automatic extraction of thumbnails. In addition, less work is available on the best method for creating memorable and meaningful thumbnails from text or PDF documents. It is important to note that even if automatic extraction is difficult for some documents, our abstract icons also can provide a significant improvement compared to standard un-augmented linear control widgets.

In future, we plan to explore the use of artificial landmarks in realistic application settings with large contents, explore the automatic extraction of meaningful landmarks, and determine the ideal number of landmarks for different document lengths. We will also carry out studies to analyze interference between landmarks and content, and explore new designs to determine how different levels of visual salience interact with the development of spatial memory.

5.3.7 Conclusions

Linear control widgets can provide a stable spatial representation of a document, and enable efficient navigation and revisitation. We explored the potential of artificial landmarks in improving spatial learning and revisitation of locations for linear documents. Linear control widgets for two applications (a media player and a PDF viewer) were augmented with artificial landmarks – either arbitrary abstract icons or thumbnails extracted from the content – to help users form a spatial understanding of the document while using the interfaces. We compared standard widgets (with no landmarks) to widgets augmented with artificial landmarks. Both artificial-landmark conditions improved performance, with the *thumbnails* condition performing best in terms of revisitation time, errors, perceived effort, and preference. Our studies show that artificial landmarks are a simple and valuable method for improving navigation in linear documents.

5.3.8 Acknowledgments

This work was supported by Natural Sciences and Engineering Research Council of Canada. Our special thanks to Miranda Miller, who recruited and ran participants.

5.4 SUMMARY OF MANUSCRIPT C

In this work, I investigated the use of artificial landmarks in linear control widgets to improve spatial memory development in linear documents, and whether landmarks can aid revisitation in those documents. I carried out a study with two types of content (a video and a PDF document) to test the effects of adding artificial landmarks in the 1D controllers. The study compared standard widgets (with no landmarks) to two landmark-augmented designs: one that placed random abstract icons in the body of the 1D widget and one that added document thumbnails to the widgets.

The study results suggest that artificial landmarks can help people develop spatial memory of linear documents, similar to command selection interfaces. Also, landmarks can support users in remembering linear document locations and improve revisitation performance. There were two main results of this study:

- Both artificial-landmark interfaces were faster than the standard widget for both media player and PDF viewer applications.
- Of the two types of artificial landmarks, the thumbnails condition was the fastest and most accurate and was strongly preferred by participants.

5.4.1 Contributions

The study provides three main contributions. First, it demonstrates two new designs to augment linear control widgets of linear document interaction interfaces with artificial landmarks. Second, this study shows that artificial landmarks can significantly improve revisitation performance in linear documents. With two different documents (a video and a PDF) and two forms of controllers (a horizontal slider and a vertical scrollbar), this work demonstrated the benefits of augmenting linear controllers with landmarks. Last, it adds to our knowledge about spatial learning and spatial retrieval in GUIs, and shows that users can successfully learn and revisit locations on individual widgets using spatial memory and landmarks.

5.4.2 Relevance in Context

In the context of my dissertation, Manuscript C broadens the investigation of artificial landmarks to a second category of interfaces – linear document revisitation interfaces equipped with linear controllers. This study extends the findings of Manuscript B, presented in Chapter 4, and introduces two new types of artificial landmarks (random icons and thumbnails from documents) that can be used in linear control widgets to help users develop spatial memory in those interfaces.

CHAPTER 6

DISCUSSION AND CONCLUSION

In this dissertation, I explored the idea of using landmarks in graphical interfaces to support users' development of spatial memory and enable users to become experts with an interface quickly. With a series of interrelated studies, my research investigated if and how landmarks in graphical user interfaces (GUIs) can help users learn and remember the locations of graphic elements (e.g., commands in ribbons and toolbars, and episodes in linear documents). The research contributes new knowledge that landmarks do help people develop memories of GUI elements and that the addition of artificial landmarks in GUIs can significantly improve the usability of interfaces by providing stable spatial reference frames for users to learn and remember locations in the interfaces quickly.

The three manuscripts in this dissertation focused on individual aspects of using landmarks in GUIs and investigated them in various contexts, such as the type of interactions (i.e., command selection and linear document revisitation), their level of realism (i.e., commercially available standard interfaces and prototype interfaces), and their coverage level (i.e., landmarks at the level of an entire interface and the level of individual widgets). Each manuscript presents a discussion of its findings (see Sections 3.3.5, 4.3.11, and 5.3.6); therefore, in this chapter, I briefly revisit the key findings from each manuscript and discuss the overall contributions in the context of existing literature. Also, I discuss lessons learned across the studies and outline the avenues my research has opened for future exploration.

6.1 SUMMARY OF THE MANUSCRIPTS

The following sections provide a brief overview of the three manuscripts presented in this dissertation. I revisit the problem statement of each paper and highlight the key findings along with their contributions.

6.1.1 Summary of Manuscript A – Landmarks in Real-World GUIs

Manuscript A addressed the problem of not knowing how people develop spatial memory of graphical commands or tools in existing graphical user interfaces and whether landmarks contribute to that spatial learning process. Although the graphical interfaces of computer applications consist of a large number of commands, users can still develop expertise with the interfaces. This manuscript presented semi-structured interviews with frequent users of four commercially available desktop applications: Microsoft Word, Facebook, Adobe Photoshop, and Adobe Reader. Findings revealed that users developed mental images of the frequently used interfaces. Further analyses of those cognitive images elicited from participants' minds disclosed how spatial memory development occurs in those four interfaces.

My findings indicated that people used four types of landmarks in the four GUIs, involving items available both inside and outside those interfaces. First, the layout of an interface containing all commands and elements of an interface served as a general reference frame that provided users with relative location information about a command. Second, a group of commands having a clearly identified boundary inside a graphical interface served as another type of landmark. Third, features of commands' icons, such as shape, size, colour, and even the meaning of icons representing commands, were consistently referred to by users during the study—indicating that these features acted as landmarks. Last, edges and corners available both inside and outside the GUI environment were stable landmarks that helped users learn and recall the location of commands.

Manuscript A filled a gap in our understanding by revealing novel and fundamental insights into the process of spatial memory development in standard GUIs. It showed that the existing design elements and features present in GUIs could act as landmarks, and people can learn and recall the locations of commands in the interfaces by relying on these landmarks. This research also confirmed one of the primary problems addressed in this dissertation by revealing that the lack of proper landmarks in some regions of GUIs made spatial learning and recall commands challenging for users.

6.1.2 Summary of Manuscript B – Testing the Effectiveness of Artificial Landmarks

Manuscript B addressed the problem that graphical command selection interfaces lack adequate landmarks to support efficient learning and remembering of the commands' locations. By carrying out three related studies with three progressively larger command-set sizes (small: 64, medium: 96, and large: 160 commands) and two novel artificial landmarks (grey coloured blocks and an image as the menu backdrop), this manuscript displayed new evidence that the use of artificial landmarks in GUIs can support efficient learning and recall of command locations. The results indicated that novices could quickly develop spatial memory of available commands in an interface and quickly became experts when landmarks were present, particularly when the interface consisted of many commands. Overall, the potential for using artificial landmarks increased proportionally with the size of a command set available in an interface.

Manuscript B makes several contributions to HCI. It provides evidence that the use of artificial landmarks in command selection GUIs can improve users' performance in learning and recalling the locations of commands. This manuscript introduced two types of artificial landmarks: abstract colour blocks as anchors and an image as the menu background. Another contribution of this work was an empirical comparison between these two landmarking strategies in three separate studies. My work demonstrated that both types supported users' spatial learning of commands. Still, even though the image offered richer landmarks than the abstract colour blocks, the simpler landmarks performed better and were preferred by participants in all three studies. This result suggests that simple landmarks may be more valuable than visually rich but complex image landmarks, a valuable takeaway message for designers and researchers focusing on developing future graphical command selection interfaces leveraging spatial memory.

6.1.3 Summary of Manuscript C – Artificial Landmarks in Linear Controllers

The problem addressed in Manuscript C was that 1D control widgets (i.e., sliders or scrollbars) in linear document interaction interfaces, such as a media player or a document reader, do not provide enough support for the efficient development of spatial memory of document locations. Banking on the findings of Manuscript B, this manuscript investigated whether the use of artificial landmarks in linear controllers can help users learn and remember the spatial locations of episodes from a linear document. Since most 1D control widgets are horizontal or vertical, exploring a horizontal slider and a vertical scrollbar was a natural choice. In order to adapt from command selection interfaces (Manuscript B) to linear widgets, Manuscript C introduced two new artificial landmarks: abstract icons and thumbnails extracted from the content of the document.

Similar to the findings with command selection interfaces, the results presented in Manuscript C indicated that using artificial landmarks in linear control widgets can help people develop spatial memory of linear documents. The results suggest that users can leverage landmarks to learn the locations of episodes from linear documents, and later, they can quickly revisit those episodes relying on the landmarks. In fact, both landmark-augmented interfaces performed faster than the standard non-landmarked widgets for both the media player and PDF viewer applications. Additionally, of the two artificial landmarks tested in Manuscript C, the thumbnail landmark was the most efficient in terms of completion time and accuracy, and participants strongly preferred it.

Manuscript B revealed new knowledge that using artificial landmarks in command selection interfaces can help people develop spatial memory of an interface. With Manuscript C, I confirmed and extended that knowledge by investigating landmarks in a different context—linear document interaction interfaces. In particular, Manuscript C tested landmarks' benefits on individual widgets (i.e., linear controllers of document viewers), while Manuscript B focused on the entire space of an interface. I also demonstrated two new artificial landmarks to augment 1D controllers of linear document interaction interfaces. First, abstract monochrome icons were used as landmarks to augment the controllers. Thumbnails extracted from the documents were the second type of landmark. These landmarks provided valuable insights into how two different artificial landmarks—one having no relation to document content other than spatial references (abstract icons) and the other displaying contents along with spatial references (thumbnails)—contribute to spatial learning and revisitation. Additionally, the study revealed that artificial landmarks could significantly improve the performance of two different types of linear documents (videos and PDFs).

6.1.4 Results in the Context of Existing Literature

Various human factors and learning theories can influence the development of users' expertise with a computer interface, as described in Chapter 2. Since this dissertation explored the use of landmarks in GUIs as a way to support users develop expertise with a GUI by enabling efficient learning and remembering of commands' locations, the findings presented in the dissertation can be better understood in the context of research on spatial memory [23,208], spatial knowledge acquisition models [195], and the stages of skill development [77,188,213].

The overall findings of the three manuscripts indicate that the use of landmarks in GUIs can help users develop spatial memory of the commands or contents of a GUI, following the model of spatial knowledge acquisition [195]. Since people first acquire knowledge of landmarks after entering a new area (i.e. landmark knowledge) [195], GUI-based landmarks are most likely to influence the acquisition of knowledge in the latter two stages (i.e., route and survey). The results from the three manuscripts support that model's validity in the context of graphical user interfaces and suggest that similar to real-life location learning, landmarks are equally valuable for spatial knowledge acquisition in graphical environments.

Second, Manuscript A explored the development of spatial memory in GUIs and sought answers to whether landmarks contribute to this spatial learning process by carrying out a study with expert users of four GUIs. Like any other skill, the skill of location learning can be developed through three sequential stages: cognitive, associative, and autonomous [77,188,213], where people in the autonomous stage are generally considered experts. Therefore, it is highly likely that investigating experts of an interface would provide the most reasonable insights into how users learned locations in the interface when they were novices (i.e., in the cognitive stage). Manuscript A revealed that the existing features and elements in GUIs, such as the layout of a GUI or a device's corners, acted as landmarks – a stable frame of reference – upon which people heavily relied to learn and remember the locations of commands in the GUIs.

Third, Manuscript A also indicated that inadequate landmarks in certain regions of GUIs could have made it difficult for users to develop a clear spatial memory of the GUI; as a result, users in the study struggled to remember commands from those areas accurately (see Chapter 3). Motivated by these findings and inspired by the benefits of landmarks in real-life location learning [15,143], my dissertation explored ways to augment GUIs with landmarks. Since spatial learning begins at the cognitive stage of skill development [77,188,213], manuscripts B and C introduced artificial landmarks in GUIs at this stage so that maximum learning benefits could be achieved. Results indicated that the use of artificial landmarks in GUIs significantly improved spatial learning and revisitation, particularly when a large number of commands were present in an interface. One reason could be that novice users learned the locations in GUIs using the landmarks; therefore, in the autonomous stage of location learning, expert users could recall the locations from their memory, leveraging the landmarks available in the GUIs.

In summary, my findings suggested that the addition of landmarks in GUIs can improve expertise development in GUIs; therefore, researchers should explore the idea of using landmarks in GUIs to improve expertise development in GUIs.

6.2 LESSONS FOR GUI LANDMARK RESEARCH

Throughout the three manuscripts presented in this dissertation, important lessons are learned concerning the development of mental images of interfaces, internal and external validity control, design of landmarks, learning curves, and interfaces' aesthetics versus landmarks' use in graphical interfaces. These lessons are not confined to individual studies. They also guide future research focusing on augmenting GUIs with landmarks and the development of spatial memory in graphical interfaces.

6.2.1 Spatial Image of an Interface

A valuable lesson from the Manuscript A of this dissertation is that users develop spatial images of interfaces they frequently use. Spatial images of interfaces are essentially mental representations of GUIs (similar to the 'cognitive maps' [139,212]) – these spatial images of GUIs are the accumulation of users' subjective perception of a GUI. Therefore, one person's spatial image of an interface can vary from another person's image of the same interface if they focus on two different aspects of the interface. However, Manuscript A revealed one similarity among these spatial images of interfaces: the presence of landmarks in GUIs and users' reliance on these landmarks to learn and later recall the location of a command in a GUI.

Since landmarks were evidently present in interfaces' spatial images and influenced users' learning and remembering performance of locations in GUIs, they could also influence the development of the spatial images. Therefore, researchers and designers should practice caution when designing GUIs, particularly while selecting and placing graphical elements in GUIs, so that people can easily develop spatial images of the GUIs. Care should also be taken to maintain the stability of a GUI because developing a spatial image of an interface requires adequate time and practice, and it will be of no use if the interface continuously changes.

6.2.2 Internal and External Validity

Internal validity and external validity are two vital concepts in the realm of experimental research [144]. While internal validity concerns the accuracy of experimental procedures so that an experiment can generate reliable results, on the other hand, external validity refers to the usefulness and generalizability of results so that they can be applicable to real-life contexts. The five studies reported in this dissertation involved both qualitative and quantitative research methods, and investigated existing and prototype interfaces. Like any other experimental research, in my studies, I also dealt with the tension between external and internal validity [144], so that the studies could generate sufficiently accurate and generalizable results.

Manuscript A investigated a fundamental question of how people develop spatial memory of interfaces by carrying out a semi-structured interview of expert users of four desktop interfaces. In order to improve the internal validity of the study, I maintained consistency while recruiting participants for each of the four interface conditions, and all participants went through the same set of questions and tasks. For achieving high external validity, however, four popular yet unique desktop applications' interfaces were chosen that covered various aspects of standard GUIs. For example, the applications varied in the number of commands (from large to small command sets), the structures (toolbars at the top to side panels), and their intended tasks (text editing, image editing, social media). The consideration of these real-life factors while selecting the study interfaces contributed to the generalizability of the results.

The subsequent two manuscripts determined if additional landmarks can be used in GUIs to improve spatial memory development for GUIs. Manuscripts B and C tested the use of artificial landmarks in command selection interfaces and linear document viewer interfaces. To accomplish the overall goal of my dissertation, I compared landmark-augmented prototype interfaces to their non-landmarked counterparts while ensuring experimental results remained internally sound and externally valid. To achieve external validity in the studies, I chose two popular interaction scenarios involving spatial memory (i.e., selecting commands from a menu and visiting episodes in linear documents) and replicated them in laboratory studies using realistic prototype interfaces. I improved the internal validity of these experiments by randomly sampling participants and counterbalancing the interface conditions.

6.2.3 Design of Landmarks

Across five studies reported in the three manuscripts of this dissertation, there are several lessons that designers could leverage while designing landmarks for GUIs. In addition to demonstrating landmarks' benefits in learning and remembering locations in GUIs, the results indicated that not all landmarks provided the same level of spatial benefits – surprisingly simple-abstract landmarks outperformed concrete and relatively feature-heavy landmarks (e.g., grey coloured blocks vs. an image as the backdrop a menu). Therefore, designers require to practice caution while designing landmarks for an interface. As our study suggested, designers can consider choosing simple and abstract landmarks over feature-rich and concrete landmarks to augment GUIs because an inappropriate choice of landmark for an interface might impede the process of spatial memory development rather than improving it.

In addition, Manuscripts B and C compared different landmarks' ability to support spatial learning and recall with multiple studies in two distinct interaction scenarios: selecting commands from menus and visiting episodes in linear documents, respectively. Findings from these studies do suggest that designers can include landmarks in GUIs. Manuscripts B and C, however, established the value of landmark augmentation in GUIs with two separate sets of landmarks (Manuscript B: abstract colour blocks and a backdrop image; Manuscript C: random icons and extracted thumbnails). Therefore, designers need to be careful before using the same landmark in different interfaces. For example, the icon landmarks used in liner controllers (in Manuscript C) were found to be beneficial in revisiting episodes from linear documents, but the icon landmarks could interfere with actual command icons if they were used in command selection interfaces.

Finally, Manuscript A identified four types of landmarks (i.e., interface layout, command groups, edges and corners, and icon visuals) present in existing interfaces and revealed their benefits in developing spatial memory of commands. However, these existing design elements and features of GUIs are not explicitly considered landmarks by GUI designers. This dissertation suggests that designers can incorporate the 'idea of landmarks' in the GUI design process and better use the newly discovered knowledge about existing design elements to help users develop spatial memory of GUIs.

6.2.4 Learning Curves

The performance of learning and remembering locations in GUIs can be described using a learning curve [4,44], because a learning curve reflects a learner's performance on a task and the amount of time spent to complete that task. Since learning begins when a user first starts to carry out a task, I primarily explored the early stages of learning while investigating ways to improve the performance of learning and remembering locations in GUIs.

Spatial interactions in graphical user interfaces are primarily comprised of two fundamental spatial actions: location learning and location retrieving. Cockburn et al.'s model for predicting menu performance [50] suggests that location learning in a graphical menu predominantly occurs through a slow visual search, particularly when users are novices because they have little knowledge of the menu. However, from those initial interactions, people start to develop an understanding of graphical elements' locations in a menu, thanks to spatial memory [195]. Also, as users progress from novices to experts, they tend to rely more on quickly recalling locations from memory instead of engaging in a slow visual search in the menu. Since people start to practice this memory-based recall in the early stages of learning [50,120], a drastic improvement in the performance of finding graphical items' locations can be visible in the early stages of a learning curve [77].

Researchers in HCI have investigated spatial interactions (i.e., learning and retrieval of locations) in GUIs and measured users' performance in carrying out spatial tasks that were repeated multiple times, usually more than 10 times [94,95,220]. Since one primary goal of this dissertation was to test the use of landmarks in GUIs to improve spatial memory, Manuscripts B and C involved users repeating spatial tasks over multiple blocks, similar to early research [94,95,220]. The three studies presented in Manuscript B, with multiple repeating blocks (varied between 10 and 18), demonstrated that most changes in the performance of spatial interactions occurred within the first few blocks. Results indicated that after a dramatic change in the performance during the early stages of learning locations (usually in the first 5 blocks), the performance reaches a plateau. Since most spatial knowledge development (i.e., spatial learning) occurred in the early blocks [50,120], studies investigating spatial performance in GUIs could lower the number of blocks. Therefore, I carried out a study presented in Manuscript C consisting of only five blocks of spatial tasks, which yielded results comparable to early studies concerning spatial memory [7,95,183]. The smaller number of blocks in a study would improve participants' comfort by minimizing the study's overall runtime and task load; however, further research is required to understand if reducing the number of blocks impacts the balance between internal and external validity of a study.

6.2.5 GUI Aesthetics vs. Landmarks in GUIs

A corresponding finding in my dissertation, particularly in Manuscript B, was that the value of landmarks increased when the number of items in an interface and the number of landmarks increased. As a result, it might be tempting to augment GUIs with many landmarks with the hope of improving the usability of a GUI; however, adding new graphical objects (i.e., artificial landmarks) to an existing interface may compromise its aesthetics. Therefore, designers need to be extra careful while adding landmarks in GUIs.

Findings from Manuscripts B and C suggest that the addition of artificial landmarks in GUIs can improve the performance of learning and remembering locations in GUIs; however, it is still unclear what can be used as landmarks in an interface? The answer is dependent on the interfaces where those spatial interactions take place. For example, in command selection interfaces like CommandMaps [183] that consist of several icons representing commands, it will be unwise to use the icon-landmarks that were used in linear document revisitation interfaces (Manuscript C). Additional icons in an interface already populated with icons would provide little to no landmark benefit unless they are visually different from their surroundings—a key characteristic of a landmark [143,210]. Similarly, although solid grey-colour blocks were proven useful in command selection interfaces (Manuscript B), they might become inefficient in the 1D controllers of linear document viewers as it would be challenging to differentiate among several similar-looking blocks in the linear controller. Therefore, it is essential to explore available features of graphical items (e.g., visual properties: shape, size, and colour; meanings represented by icon visuals; distinctiveness) in order to identify suitable landmarks for interfaces.

Another question relevant to the use of landmarks in GUIs and maintaining the aesthetics of GUIs that remained unanswered is how many landmarks are suitable for a GUI and where those landmarks should be included. This dissertation indicates that there are cases where the addition of landmarks may not provide any additional benefit, for example, interfaces with a small number of commands (Study 1 of Manuscript B). For a relatively smaller set of commands, usually 64 or less, I found that people's spatial memory without additional landmarks was sufficient for learning and remembering commands. However, the benefits of landmarks in spatial memory development of GUIs became evident when command set size and the number of landmarks increased (96 commands: 8 landmarks; 160 commands: 12 landmarks). In addition, Manuscript A revealed that existing design elements, features, and structures could act as landmarks that users leverage to develop spatial memory of commands in standard GUIs. The findings, therefore, encourage designers to be vigilant while including additional landmarks in interfaces. Unnecessary use of landmarks in GUIs where landmarks are already present or overuse of landmarks in a GUI can overwhelm users and disrupt the aesthetics of an interface.

In order to provide better spatial learning and remembering facilities in GUIs, this dissertation has demonstrated how existing design elements and four different artificial landmark strategies can be useful in GUIs (please see Manuscripts B and C for details). Though this work has revealed the potential for using landmarks in graphical environments to improve interfaces' usability and provided several design implications, adequate research is required to completely understand landmarks' use in GUIs without jeopardizing their aesthetics.

6.3 LIMITATIONS AND FUTURE RESEARCH DIRECTIONS

The research presented in this dissertation consisted of five different studies involving real and prototype interfaces, and there are limitations associated with each of these studies. Manuscripts A, B, and C have already discussed their respective limitations at the end of each manuscript and presented possible ways to overcome them. Therefore, in this section, I briefly present the limitations common across the studies and ways to overcome them without repeating the topics from Chapters 3 to 5. In addition, I shed light on research directions that my dissertation research can be extended in future.

6.3.1 Consideration of Multiple Platforms

In order to identify the potential of landmarks to develop spatial memory in graphical interfaces throughout the three manuscripts, I limited my focus in this dissertation to desktop interfaces only, primarily for two reasons. First, desktops are one of the most popular platforms where graphical user interfaces are predominantly seen. Additionally, desktop GUIs have been studied extensively, which allowed me to compare my findings with established literature. Second, I considered desktops as a starting context for research. Surely, including multiple platforms and devices in investigating spatial memory development in GUIs would be beneficial. However, at this early stage of this research, limiting the investigation to one platform enabled me to keep the interference from other platforms low and improve the reliability of the findings.

However, there are several other platforms currently available and becoming popular. For example, plenty of multitouch devices varies in size and shape. The 3D immersive technologies such as AR/VR have been redefining how users can interact with GUIs. Since interfaces are no longer confined to 2D spaces only, it would be interesting to see if the findings of this dissertation hold true in these new contexts. Also, multi-touch and AR/VR interfaces allow users to interact with natural objects as well as digital ones. The landmarks found useful in desktop contexts may not provide a similar level of assistance in learning and remembering locations in other contexts. Therefore, one avenue for my future research would involve investigating what can be used as landmarks in these new platforms.

Another way my research could be expanded in future is by exploring ways to design landmarks that can be used across multiple platforms and interfaces. In this dissertation, I have introduced and tested two sets of landmarks in two different interfaces: colour blocks and a backdrop image in a command selection interface (Manuscript B), random icons and thumbnails in linear controllers (Manuscript C). Although my research has established the potential for using landmarks in GUIs, it is unclear whether the same landmarks can provide similar spatial benefits in different interfaces. For example, icon landmarks used in linear controllers maybe become ineffective in command selection interfaces as they would be difficult to differentiate from existing command icons. Therefore, further research is required to design generalized landmarks that can be used in all GUI platforms, irrespective of their differences.

In addition, several interfaces allow users to customize the interface (i.e., rearrange the commands) based on their requirements. Even some interfaces continuously update the contents displayed on screens (e.g., the interface of Canvas - an online learning management system), which might create a unique challenge for users to develop spatial memory of the interface. This dissertation consciously chose to exclude these highly customizable GUIs from the analyses, primarily to keep interference from additional variables limited. However, in the future, I plan to investigate the effects of using landmarks in these customizable interfaces, particularly how and if users develop spatial memory in these GUIs.

6.3.2 Task-Centric Learning in Computers

Besides learning the locations of commands in an interface, expertise development with an application involves discovering the functionalities of commands and learning how to use those commands to carry out tasks. Since people primarily use computer applications to carry out various tasks, research indicated that learning in computers is often governed by the context of tasks people intend to perform [122,146,176]. As a result, people often end up learning only the commands related to specific tasks instead of all the commands present in an interface. Since the primary goal of this dissertation was to facilitate learning and recalling the locations of commands, as an initial step, I limited my focus only to learning the locations of commands without associating commands to any tasks. Still, I believe the findings of spatial memory and landmarks from this dissertation can remain valid in the context of task-centric learning in computers; this is because learning the locations of commands is a fundamental step in both task-centric and exploratory learning approaches [176]. However, further research is required to better understand the effects of landmarks and spatial memory in various contexts of learning computer applications. This work could be extended in the future by investigating the impact of landmarks and spatial memory by comparing various ways people learn interfaces in real life (i.e., task-centric and exploratory learning methods).

6.3.3 A Framework for Landmarks in GUIs

My dissertation encourages designers and researchers to augment GUIs with landmarks that can improve the usability of computer interfaces by making it easier for people to learn and remember locations in GUIs. However, there are open questions about designing effective landmarks for a GUI: (1) what can be used as landmarks in GUIs and (2) how many landmarks are sufficient for an interface? Some aspects of these questions were already discussed in Sections 6.2.3 and 6.2.5. In addition to discussing valuable findings of five studies, the three manuscripts provided several guidelines for designers of future interfaces. While those design guidelines and results are useful and can be considered a significant milestone towards improving our understanding of landmarks and spatial memory development in GUIs, given the large number of variables associated with designing landmarks, such as shapes, sizes, colours, and even meanings of items (Manuscript A discussed these in detail), and considering the plethora of interactive devices people currently use, it is imperative to develop a framework for landmarks in GUIs, so that designers can easily select a suitable landmark for an interface.

While developing a framework for landmarks in GUIs, three main factors can be considered, which may form the three dimensions of the framework. First, the cognitive factors associated with designing landmarks for GUIs can be comprised of human abilities to perceive landmarks, aspects of spatial memory such as learning locations, storing learned locations for short and long terms, remembering locations, and even factors related to the decay of memory. Second, the characteristics of landmarks can be another dimension of the new framework. The attributes (see Section 2.3.2.2) that turn an object into a landmark in real-life, such as visibility, and distinctiveness, are likely to be useful in graphical environments. Also, additional features such as the meaning of the icons (Manuscript A), sounds, or even natural objects (e.g., users' own hands [217,220]) can become landmarks in GUIs. Last, the factors associated with interfaces can be a dimension of the framework. Factors such as the platform where a graphical interface is going to be displayed (desktops, AR/VR, or multi-touch), the size of the interface, the number of items present in the interface, or even input methods (e.g., direct or indirect methods) can influence the performance of a landmark in a GUI.

6.3.4 Spatial Learning Retention and Transfer

Early research in Psychology suggests that one way to measure learning is through testing the recall performance immediately after learning or through retention and transfer tests, which are often carried out after 24 and 48 hours delay [51]. Since this dissertation involved measuring the performance of learning locations in GUIs, I tested users' spatial learning through immediate testing and retention, particularly in Manuscripts B and C. Although it is preferred to run retention and transfer tests with a sufficient gap between learning and testing (e.g., 24 and 48 hours gaps between learning and testing [51]), it is common in HCI research to test learning performance immediately after learning [45,81,95,126,133]. Manuscripts B and C also followed a performance testing model similar to those used in early studies. Still, the absence of testing the retention performance of spatial learning after a gap in these studies can be considered a limitation.

Although Manuscript A's study investigated users' spatial knowledge of GUIs and landmarks' use by asking people to recall locations in GUIs that they previously learnt, it did not test the effect of adding landmarks in GUIs as the studies in Manuscripts B and C did. Therefore, one way my research could be extended in the future is by testing the retention and transfer of spatial knowledge with landmarks in GUIs. My analysis indicated that landmarks could help people develop spatial memory of GUIs quickly; however, we do not know how long this memory will last. Since research suggests that people's memory decay over time [71], there is a possibility that people may forget the locations in GUIs after learning. Therefore, an interesting area of future research could explore if the use of landmarks in GUIs reduces that memory decay.

Besides, findings from Manuscript A suggest that four types of landmarks were commonly present across four different GUIs that varied in layouts and command numbers. Since the landmarks people relied on to develop the memory of commands in those GUIs were similar, it could be possible that the spatial knowledge acquired in one interface can be useful in another interface. For example, the controls for 'Close,' 'Minimize,' and 'Maximize' in a GUI typically appear at the top right corner in a Windows OS running system. This spatial knowledge can be useful in locating those controls in other applications' GUIs. Therefore, it would be worth exploring whether spatial knowledge developed in an interface can be useful or transferred in another interface with similar landmarks.

6.4 CONTRIBUTIONS

My dissertation makes several contributions to HCI and Computer Science by expanding our knowledge of spatial memory development with GUIs.

- This dissertation reveals valuable new information that people develop *cognitive images* of graphical interfaces that they regularly use (Manuscript A). Users who are new to an interface start interacting with it through visually locating commands. These interactions help users develop an image of that interface in their minds. My dissertation also provides evidence that people, particularly experienced users, rely on these spatial images to recall the locations of commands from memory.

- My dissertation introduces and demonstrates a method to elicit spatial images of interfaces from users' minds (Manuscript A). It also analyses the mental images of four commercially available desktop applications' interfaces: Microsoft Word, Facebook, Adobe Photoshop, and Adobe Reader. Manuscript A revealed novel information about spatial memory development in GUIs: standard interfaces' existing features and structural elements can act as reliable landmarks that can be grouped into four categories. It also revealed that people strongly rely on these four types of landmarks to learn and recall the locations of commands in an interface.

- With four studies presented in Manuscripts B and C, this dissertation empirically demonstrates for the first time how *artificial landmarks* can be used in GUIs to aid in better location learning of graphical elements (e.g., commands/tools and episodes in linear documents such as videos and PDFs). It also reveals that the presence of artificial landmarks in GUIs can enable users' efficient revisitation to spatial locations in two different contexts: command selection interfaces and linear document viewers.

- My dissertation introduces four types of artificial landmarks (abstract blocks, an image as a menu-backdrop, random icons, and thumbnails from documents) and demonstrates how these landmarks can be used to augment two prototype interfaces: command selection interfaces and linear document viewers (Manuscripts B and C). It also empirically validates how different landmarks perform in interfaces that vary in the number of commands and types. Overall, all landmarks provided spatial benefits; however, simple abstract blocks and thumbnail landmarks were proven more beneficial in command selection GUIs and linear document viewers.

- Finally, this dissertation provides guidelines for designers of future interfaces (all manuscripts) so that they can design more memorable GUIs with the help of landmarks, be it by conscious use of existing design elements or by introducing digitally crafted artificial landmarks.

6.5 CONCLUSION

Two-dimensional graphical user interfaces present graphical items (e.g., commands) at particular locations in the interfaces that users require to find and visit in order to carry out tasks on computers. Users can efficiently complete tasks on computers by learning and remembering those locations in GUIs, since recalling locations from memory is faster than visually searching for them. Spatial memory enables people to learn and remember locations in an area. However, learning and remembering locations in GUIs can be complicated and show as these interfaces lack adequate landmarks. In order to understand clearly how spatial memory development occurs in GUIs and find ways to assist users in efficient location learning and recall, I carried out five studies exploring the use of landmarks in GUIs. The first study investigated interfaces of four standard desktop applications: Microsoft Word, Facebook, Adobe Photoshop, and Adobe Reader, and the other four studies tested the use of 'artificial landmarks' in two prototype GUIs (command selection interfaces and linear document viewers) against respective non-landmarked versions. Results revealed that existing features and design elements in standard GUIs can act as reliable landmarks in supporting learning and remembering locations in GUIs, and that artificially created landmarks can significantly improve spatial memory development and support rapid user expertise development in GUIs.

This dissertation makes several contributions by providing new knowledge to the field of Computer Science. First, people develop mental images of GUIs they frequently use, and they heavily rely on four types of landmarks (interface layout, command groups, corner and edges, and icon visuals) to learn and recall the locations of commands in standard GUIs. Second, artificial landmarks can significantly improve spatial learning performance, particularly when the number of commands increases, and simple abstract landmarks can outperform feature-rich concrete landmarks in providing spatial benefits. Third, it provides guidelines for designers so that they can design improved GUIs augmented with landmarks. Overall, this dissertation demonstrates that landmarks can be a valuable addition to GUIs to improve graphical interfaces' memorability and usability by enabling people to learn and remember locations in GUIs quickly.

REFERENCES

[1] David Abrams, Ron Baecker, and Mark Chignell. 1998. Information archiving with bookmarks. In *Proceedings of the SIGCHI conference on Human factors in computing systems - CHI '98*, ACM Press, New York, NY, USA, 41–48. DOI:https://doi.org/10.1145/274644.274651

[2] Eytan Adar, Jaime Teevan, and Susan T. T. Dumais. 2008. Large scale analysis of web revisitation patterns. In *Proceeding of the SIGCHI conference on Human factors in computing systems - CHI '08*, ACM, New York, NY, USA, 1197–1206. DOI:https://doi.org/10.1145/1357054.1357241

[3] Christopher Ahlberg and Ben Shneiderman. 1994. The alphaslider: a compact and rapid selector. In *Proceedings of the SIGCHI conference on Human factors in computing systems - CHI '94*, ACM, New York, NY, USA, 226. DOI:https://doi.org/10.1145/259963.260406

[4] David Ahlström, Andy Cockburn, Carl Gutwin, and Pourang Irani. 2010. Why it's quick to be square: modelling new and existing hierarchical menu designs. In *Proceedings of the 28th international conference on Human factors in computing systems - CHI '10*, ACM Press, New York, New York, USA, New York, USA, 1371–1380. DOI:https://doi.org/10.1145/1753326.1753534

[5] Abir Al-Hajri, Gregor Miller, Matthew Fong, and Sidney S. Fels. 2014. Visualization of personal history for video navigation. In *Proceedings of the SIGCHI conference on Human factors in computing systems - CHI '14*, ACM, New York, NY, USA, 1187–1196. DOI:https://doi.org/10.1145/2556288.2557106

[6] Gerianne M. Alexander, Mark G. Packard, and Bradley S. Peterson. 2002. Sex and spatial position effects on object location memory following intentional learning of object identities. *Neuropsychologia* 40, 8 (2002), 1516–1522. DOI:https://doi.org/10.1016/S0028-3932(01)00215-9

[7] Jason Alexander, Andy Cockburn, Stephen Fitchett, Carl Gutwin, and Saul Greenberg. 2009. Revisiting read wear: analysis, design, and evaluation of a footprints scrollbar. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems - CHI '09*, ACM, New York, NY, USA, 1665–1674. DOI:https://doi.org/10.1145/1518701.1518957

[8] Gary L. Allen. 1981. A developmental perspective on the effects of "subdividing" macrospatial experience. *Journal of Experimental Psychology: Human Learning and Memory* 7, 2 (March 1981), 120–132. DOI:https://doi.org/10.1037/0278-7393.7.2.120

[9] Gary L. Allen, Alexander W. Siegel, and Richard R. Rosinski. 1978. The role of perceptual context in structuring spatial knowledge. *Journal of Experimental Psychology: Human*

Learning & Memory 4, 6 (1978), 617–630. DOI:https://doi.org/10.1037/0278-7393.4.6.617

[10] Amir Amedi, Rafael Malach, Talma Hendler, Sharon Peled, and Ehud Zohary. 2001. Visuo-haptic object-related activation in the ventral visual pathway. *Nature Neuroscience* 4, 3 (2001), 324–330. DOI:https://doi.org/10.1038/85201

[11] Fraser Anderson and Walter F. Bischof. 2013. Learning and performance with gesture guides. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems - CHI '13*, ACM Press, New York, New York, USA, 1109–1118. DOI:https://doi.org/10.1145/2470654.2466143

[12] John R. Anderson. 1982. Acquisition of cognitive skill. *Psychological Review* 89, 4 (July 1982), 369–406. DOI:https://doi.org/10.1037/0033-295X.89.4.369

[13] John R. Anderson and Gordon H. Bower. 1980. *Human Associative Memory*. Psychology Press. DOI:https://doi.org/10.4324/9780203780831

[14] John Robert Anderson. 2000. *Learning and memory: An integrated approach, 2nd ed.* John Wiley & Sons Inc, Hoboken, NJ, US.

[15] Jackie Andrade and Peter Meudell. 1993. Short report: is spatial information encoded automatically in memory? *The Quarterly journal of experimental psychology. A, Human experimental psychology* 46, 2 (May 1993), 365–375. DOI:https://doi.org/10.1080/14640749308401051

[16] Donald Appleyard. 1969. Why Buildings Are Known. *Environment and Behavior* 1, 2 (December 1969), 131–156. DOI:https://doi.org/10.1177/001391656900100202

[17] Wendy Ark, D Christopher Dryer, Ted Selker, and Shumin Zhai. 1998. Landmarks to aid navigation in a graphical user interface. In *Workshop on Personalised and Social Navigation in Information Space*, Stockholm. Retrieved September 15, 2016 from https://www.researchgate.net/publication/235328897_Landmarks_to_aid_navigation_in_a_graphical_user_interface

[18] Marieke van Asselen, Roy P. C. Kessels, L. Jaap Kappelle, and Albert Postma. 2008. Categorical and coordinate spatial representations within object-location memory. *Cortex* 44, 3 (March 2008), 249–256. DOI:https://doi.org/10.1016/j.cortex.2006.05.005

[19] Richard C. Atkinson and Richard M. Shiffrin. 1968. Human Memory: A Proposed System and its Control Processes. In *Psychology of Learning and Motivation - Advances in Research and Theory*. Academic Press, 89–195. DOI:https://doi.org/10.1016/S0079-7421(08)60422-3

[20] Emanuel Averbach and Abner S. Coriell. 1961. Short-Term Memory in Vision. *Bell System Technical Journal* 40, 1 (January 1961), 309–328. DOI:https://doi.org/10.1002/j.1538-7305.1961.tb03987.x

[21] Alan Baddeley. 1992. Working memory. *Science* 255, 5044 (1992), 556–559. DOI:https://doi.org/10.1126/science.1736359

[22] Alan Baddeley. 2000. The episodic buffer: a new component of working memory? *Trends in Cognitive Sciences* 4, 11 (November 2000), 417–423. DOI:https://doi.org/10.1016/S1364-6613(00)01538-2

[23] Alan D. Baddeley. 1990. *Human memory : theory and practice*. Hove, [England] : Lawrence Erlbaum.

[24] Alan D. Baddeley. 1990. The development of the concept of working memory: implications and contributions of neuropsychology. In *Neuropsychological Impairments of Short-Term Memory*. Cambridge University Press, 54–73. DOI:https://doi.org/10.1017/CBO9780511665547.004

[25] Alan D. Baddeley. 1999. *Essentials of human memory*. Psychology Press, Hove England.

[26] Alan D. Baddeley. 2004. The Psychology of Memory. In *The Essential Handbook of Memory Disorders for Clinicians*. John Wiley & Sons, Ltd., 1–14.

[27] Alan D. Baddeley and Graham Hitch. 1974. Working Memory. In *Psychology of Learning and Motivation - Advances in Research and Theory*. Academic Press, 47–89. DOI:https://doi.org/10.1016/S0079-7421(08)60452-1

[28] Alan D. Baddeley and Robert H. Logie. 1999. Working Memory: The Multiple-Component Model. In *Models of Working Memory*. Cambridge University Press, 28–61. DOI:https://doi.org/10.1017/CBO9781139174909.005

[29] Gilles Bailly, Eric Lecolinet, and Laurence Nigay. 2008. Flower menus: a new type of marking menu with large menu breadth, within groups and efficient expert mode memorization. In *Proceedings of the working conference on Advanced visual interfaces - AVI '08*, ACM Press, New York, New York, USA, 15–22. DOI:https://doi.org/10.1145/1385569.1385575

[30] Olivier Bau and Wendy E. Mackay. 2008. OctoPocus: a dynamic guide for learning gesture-based command sets. In *Proceedings of the ACM symposium on User interface software and technology - UIST '08*, ACM Press, New York, New York, USA, 37–46. DOI:https://doi.org/10.1145/1449715.1449724

[31] Patrick Baudisch and Ruth Rosenholtz. 2003. Halo: a technique for visualizing off-screen objects. In *Proceedings of the conference on Human factors in computing systems - CHI '03*, ACM Press, New York, New York, USA, 481–488. DOI:https://doi.org/10.1145/642611.642695

[32] Frank Bentley and Janet Murray. 2016. Understanding Video Rewatching Experiences. In *Proceedings of the ACM International Conference on Interactive Experiences for TV and*

Online Video - TVX '16, ACM Press, New York, NY, USA, 69–75. DOI:https://doi.org/10.1145/2932206.2932213

[33] Meera Blattner, Denise Sumikawa, and Robert Greenberg. 1989. Earcons and Icons: Their Structure and Common Design Principles. *Human-Computer Interaction* 4, 1 (March 1989), 11–44. DOI:https://doi.org/10.1207/s15327051hci0401_1

[34] James C. Bliss, Hewitt D. Crane, Phyllis K. Mansfield, and James T. Townsend. 1966. Information available In brief tactile presentations. *Perception & Psychophysics* 1, 4 (July 1966), 273–283. DOI:https://doi.org/10.3758/BF03207391

[35] John Boreczky, Andreas Girgensohn, Gene Golovchinsky, and Shingo Uchihashi. 2000. An interactive comic book presentation for exploring video. In *Proceedings of the SIGCHI conference on Human factors in computing systems - CHI '00*, ACM Press, New York, NY, USA, 185–192. DOI:https://doi.org/10.1145/332040.332428

[36] Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qualitative Research in Psychology* 3, 2 (January 2006), 77–101. DOI:https://doi.org/10.1191/1478088706qp063oa

[37] Virginia Braun and Victoria Clarke. 2019. Reflecting on reflexive thematic analysis. *Qualitative Research in Sport, Exercise and Health 11*, 589–597. DOI:https://doi.org/10.1080/2159676X.2019.1628806

[38] James R. Brockmole, Monica S. Castelhano, and John M. Henderson. 2006. Contextual cueing in naturalistic scenes: Global and local contexts. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 32, 4 (July 2006), 699–706. DOI:https://doi.org/10.1037/0278-7393.32.4.699

[39] James R. Brockmole and John M. Henderson. 2006. Using real-world scenes as contextual cues for search. *Visual Cognition* 13, 1 (January 2006), 99–108. DOI:https://doi.org/10.1080/13506280500165188

[40] Lee R. Brooks. 1968. Spatial and verbal components of the act of recall. *Canadian Journal of Psychology/Revue canadienne de psychologie* 22, 5 (1968), 349–368. DOI:https://doi.org/10.1037/h0082775

[41] Donald Byrd. 1999. A scrollbar-based visualization for document navigation. In *Proceedings of the fourth ACM conference on Digital libraries - DL '99*, ACM Press, New York, NY, USA, 122–129. DOI:https://doi.org/10.1145/313238.313283

[42] John R. Callahan, Don Hopkins, Mark D. Weiser, and Ben Shneiderman. 1988. An empirical comparison of pie vs. linear menus. In *Proceedings of the SIGCHI conference on Human factors in computing systems - CHI '88*, ACM Press, New York, New York, USA, 95–100. DOI:https://doi.org/10.1145/57167.57182

[43] Eduardo Camina and Francisco Güell. 2017. The Neuroanatomical, Neurophysiological and Psychological Basis of Memory: Current Models and Their Origins. *Frontiers in Pharmacology* 8, JUN (June 2017), 438. DOI:https://doi.org/10.3389/fphar.2017.00438

[44] Stuart K. Card, Thomas P. Moran, and Allen Newell. 1983. *The Psychology of Human-Computer Interaction*. Lawrence Erlbaum Associates, Inc., Hillsdale, NJ.

[45] Febi Chajadi, Md. Sami Uddin, and Carl Gutwin. 2020. Effects of Visual Distinctiveness on Learning and Retrieval in Icon Toolbars. In *Proceedings of the 46th Graphics Interface Conference, GI 2020*, Toronto, ON, Canada, 11.

[46] William G. Chase. 1986. Visual information processing. In *Handbook of perception and human performance, Vol. 2: Cognitive processes and performance.* John Wiley & Sons, Oxford, England, England, 1–71.

[47] Ling Chen, Gen-Cai Chen, Cheng-Zhe Xu, Jack March, and Steve Benford. 2008. EmoPlayer: A media player for video clips with affective annotations. *Interacting with Computers* 20, 1 (January 2008), 17–28. DOI:https://doi.org/10.1016/j.intcom.2007.06.003

[48] Andy Cockburn and Carl Gutwin. 2009. A Predictive Model of Human Performance With Scrolling and Hierarchical Lists. *Human-Computer Interaction* 24, 3 (July 2009), 273–314. DOI:https://doi.org/10.1080/07370020902990402

[49] Andy Cockburn, Carl Gutwin, and Jason Alexander. 2006. Faster document navigation with space-filling thumbnails. In *Proceedings of the SIGCHI conference on Human Factors in computing systems - CHI '06*, ACM, New York, NY, USA, 1–10. DOI:https://doi.org/10.1145/1124772.1124774

[50] Andy Cockburn, Carl Gutwin, and Saul Greenberg. 2007. A predictive model of menu performance. In *Proceedings of the SIGCHI conference on Human factors in computing systems - CHI '07*, ACM Press, New York, NY, USA, 627–636. DOI:https://doi.org/10.1145/1240624.1240723

[51] Andy Cockburn, Carl Gutwin, Joey Scarr, and Sylvain Malacria. 2014. Supporting Novice to Expert Transitions in User Interfaces. *ACM Computing Surveys* 47, 2 (November 2014), 1–36. DOI:https://doi.org/10.1145/2659796

[52] Andy Cockburn, Per Ola Kristensson, Jason Alexander, and Shumin Zhai. 2007. Hard lessons: effort-inducing interfaces benefit spatial learning. In *Proceedings of the SIGCHI conference on Human factors in computing systems - CHI '07*, ACM Press, New York, New York, USA, New York, USA, 1571–1580. DOI:https://doi.org/10.1145/1240624.1240863

[53] Andy Cockburn and Bruce Mckenzie. 2001. What do web users do? An empirical analysis of web use. *International Journal of Human-Computer Studies* 54, 6 (June 2001), 903–922. DOI:https://doi.org/10.1006/ijhc.2001.0459

[54] Andy Cockburn and Bruce McKenzie. 2002. Evaluating the effectiveness of spatial memory in 2D and 3D physical and virtual environments. In *Conference on Human Factors in Computing Systems - Proceedings*, Association for Computing Machinery (ACM), New York, New York, USA, 203–210. DOI:https://doi.org/10.1145/503376.503413

[55] Andrew Coleman. 2006. *Dictionary of Psychology* (2nd ed.). Oxford University Press.

[56] Max Coltheart. 1983. Ecological necessity of iconic memory. *Behavioral and Brain Sciences* 6, 1 (1983), 17–18. DOI:https://doi.org/10.1017/S0140525X00014357

[57] Juliet Corbin and Anselm Strauss. 2014. *Basics of qualitative research: Techniques and procedures for developing grounded theory* (4th ed.). Sage publications. Retrieved from https://us.sagepub.com/en-us/nam/basics-of-qualitative-research/book235578

[58] Helen M. Couclelis, Reginald G. Golledge, Nathan D. Gale, and Waldo R. Tobler. 1987. Exploring the anchor-point hypothesis of spatial cognition. *Journal of Environmental Psychology* 7, 2 (June 1987), 99–122. DOI:https://doi.org/10.1016/S0272-4944(87)80020-8

[59] Nelson Cowan. 1999. An Embedded-Processes Model of Working Memory. In *Models of Working Memory*. Cambridge University Press, 62–101. DOI:https://doi.org/10.1017/CBO9781139174909.006

[60] Nelson Cowan, Wemara Lichty, and Tim R. Grove. 1990. Properties of Memory for Unattended Spoken Syllables. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 16, 2 (1990), 258–269. DOI:https://doi.org/10.1037/0278-7393.16.2.258

[61] Fergus I.M. Craik and Robert S. Lockhart. 1972. Levels of processing: A framework for memory research. *Journal of Verbal Learning and Verbal Behavior* 11, 6 (December 1972), 671–684. DOI:https://doi.org/10.1016/S0022-5371(72)80001-X

[62] Edward R. F. W. Crossman. 1959. A Theory of the Acquisition of Speed-Skill. *Ergonomics* 2, 2 (February 1959), 153–166. DOI:https://doi.org/10.1080/00140135908930419

[63] Mary Czerwinski and George Robertson. 1999. The Contribution of Thumbnail Image, Mouse-over Text and Spatial Location Memory to Web Page Retrieval in 3D. In *Proceedings of Seventh International Conference on Human-Computer Interaction - INTERACT '99*, IOS Press, Edinburgh, Scotland, Scotland, 163–170.

[64] Rudolph P. Darken and John L. Sibert. 1996. Wayfinding strategies and behaviors in large virtual worlds. In *Proceedings of the SIGCHI conference on Human factors in computing systems common ground - CHI '96*, ACM Press, New York, New York, USA, 142–149. DOI:https://doi.org/10.1145/238386.238459

[65] Rudy P. Darken and John L. Sibert. 1993. A toolset for navigation in virtual environments. In *UIST 1993 - Proceedings of the 6th Annual ACM Symposium on User Interface Software*

and Technology, Association for Computing Machinery, Inc, 157–165. DOI:https://doi.org/10.1145/168642.168658

[66] James Deese and Roger A Kaufman. 1957. Serial effects in recall of unorganized and sequentially organized verbal material. *Journal of experimental psychology* 54, 3 (September 1957), 180–7.

[67] Robert DeLine, Mary Czerwinski, Brian R. Meyers, Gina Venolia, Steven Drucker, and George Robertson. 2006. Code Thumbnails: Using Spatial Memory to Navigate Source Code. In *Visual Languages and Human-Centric Computing (VL/HCC'06)*, IEEE, 11–18. DOI:https://doi.org/10.1109/VLHCC.2006.14

[68] Bradford C. Dickerson and Howard Eichenbaum. 2010. The Episodic Memory System: Neurocircuitry and Disorders. *Neuropsychopharmacology* 35, 1 (January 2010), 86–104. DOI:https://doi.org/10.1038/npp.2009.126

[69] Vaibhav A. Diwadkar and Timothy P. McNamara. 1997. Viewpoint Dependence in Scene Recognition. *Psychological Science* 8, 4 (July 1997), 302–307. DOI:https://doi.org/10.1111/j.1467-9280.1997.tb00442.x

[70] Randolph D. Easton, Anthony J. Greene, and Kavitha Srinivas. 1997. Transfer between vision and haptics: Memory for 2-D patterns and 3-D objects. *Psychonomic Bulletin and Review* 4, 3 (1997), 403–410. DOI:https://doi.org/10.3758/BF03210801

[71] Hermann Ebbinghaus. 1913. *Memory: A Contribution to Experimental Psychology* (Number 3 ed.). Teachers College, Columbia University.

[72] Brian D. Ehret. 2002. Learning where to look: location learning in graphical user interfaces. In *Proceedings of the SIGCHI conference on Human factors in computing systems - CHI '02*, ACM, New York, NY, USA, 211–218. DOI:https://doi.org/10.1145/503376.503414

[73] Arne D. Ekstrom, Aiden E. G. F. Arnold, and Giuseppe Iaria. 2014. A critical review of the allocentric spatial representation and its neural underpinnings: toward a network-based perspective. *Frontiers in Human Neuroscience* 8, OCT (October 2014), 803. DOI:https://doi.org/10.3389/fnhum.2014.00803

[74] Gary W. Evans. 1980. Environmental cognition. *Psychological Bulletin* 88, 2 (September 1980), 259–287. DOI:https://doi.org/10.1037/0033-2909.88.2.259

[75] Nathan J. Evans, Scott D. Brown, Douglas J.K. Mewhort, and Andrew Heathcote. 2018. Refining the law of practice. *Psychological Review* 125, 4 (July 2018), 592–605. DOI:https://doi.org/10.1037/rev0000105

[76] Leah Findlater and Joanna McGrenere. 2004. A comparison of static, adaptive, and adaptable menus. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems - CHI '04*, ACM, New York, NY, USA, 89–96.

DOI:https://doi.org/10.1145/985692.985704

[77] Paul M. Fitts and Michael I. Posner. 1967. *Human performance*. Brooks/Cole, Oxford, England.

[78] Krzysztof Z. Gajos, Mary Czerwinski, Desney S. Tan, and Daniel S. Weld. 2006. Exploring the design space for adaptive graphical user interfaces. In *Proceedings of the working conference on Advanced visual interfaces - AVI '06*, ACM Press, New York, New York, USA, 201–208. DOI:https://doi.org/10.1145/1133265.1133306

[79] BoYu Gao, Byungmoon Kim, Jee-In Kim, and HyungSeok Kim. 2019. Amphitheater Layout with Egocentric Distance-Based Item Sizing and Landmarks for Browsing in Virtual Reality. *International Journal of Human–Computer Interaction* 35, 10 (June 2019), 831–845. DOI:https://doi.org/10.1080/10447318.2018.1498654

[80] BoYu Gao, HyungSeok Kim, Byungmoon Kim, and Jee-In Kim. 2018. Artificial Landmarks to Facilitate Spatial Learning and Recalling for Curved Visual Wall Layout in Virtual Reality. In *2018 IEEE International Conference on Big Data and Smart Computing (BigComp)*, IEEE, 475–482. DOI:https://doi.org/10.1109/BigComp.2018.00076

[81] Varun Gaur, Md. Sami Uddin, and Carl Gutwin. 2018. Multiplexing spatial memory: increasing the capacity of FastTap menus with multiple tabs. In *Proceedings of the 20th International Conference on Human-Computer Interaction with Mobile Devices and Services - MobileHCI '18*, ACM Press, New York, New York, USA, 1–13. DOI:https://doi.org/10.1145/3229434.3229482

[82] Emilien Ghomi, Stéphane Huot, Olivier Bau, Michel Beaudouin-Lafon, and Wendy E. Mackay. 2013. Arpège: learning multitouch chord gestures vocabularies. In *Proceedings of the ACM conference on Interactive tabletops and surfaces - ITS '13*, ACM Press, New York, New York, USA, 209–218. DOI:https://doi.org/10.1145/2512349.2512795

[83] Reginald G. Golledge. 1999. Human wayfinding and cognitive maps. *Wayfinding behavior: Cognitive mapping and other spatial processes* (1999), 5–45.

[84] Reginald G. Golledge, Terence R. Smith, James W. Pellegrino, Sally Doherty, and Sandra P. Marshall. 1985. A conceptual model and empirical analysis of children's acquisition of spatial knowledge. *Journal of Environmental Psychology* 5, 2 (June 1985), 125–152. DOI:https://doi.org/10.1016/S0272-4944(85)80014-1

[85] Adrienne M. Grech, Jay Patrick Nakamura, and Rachel Anne Hill. 2018. The Importance of Distinguishing Allocentric and Egocentric Search Strategies in Rodent Hippocampal-Dependent Spatial Memory Paradigms: Getting More Out of Your Data. In *The Hippocampus - Plasticity and Functions*. InTech, 105–126. DOI:https://doi.org/10.5772/intechopen.76603

[86] Saul Greenberg. 1993. *The computer user as toolsmith : the use, reuse, and organization of*

computer-based tools. Cambridge University Press.

[87] Saul Greenberg and Ian H. H. Witten. 1993. Supporting command reuse: empirical foundations and principles. *International Journal of Man-Machine Studies* 39, 3 (September 1993), 353–390. DOI:https://doi.org/10.1006/imms.1993.1065

[88] Saul Greenberg and Ian H. H. Witten. 1993. Supporting command reuse: mechanisms for reuse. *International Journal of Man-Machine Studies* 39, 3 (September 1993), 391–425. DOI:https://doi.org/10.1006/imms.1993.1066

[89] Greg Guest, Arwen Bunce, and Laura Johnson. 2006. How Many Interviews Are Enough? An Experiment with Data Saturation and Variability. *Field Methods* 18, 1 (February 2006), 59–82. DOI:https://doi.org/10.1177/1525822X05279903

[90] Sean Gustafson, Patrick Baudisch, Carl Gutwin, and Pourang Irani. 2008. Wedge: clutter-free visualization of off-screen locations. In *Proceeding of the twenty-sixth annual CHI conference on Human factors in computing systems - CHI '08*, ACM Press, New York, New York, USA, 787–796. DOI:https://doi.org/10.1145/1357054.1357179

[91] Sean Gustafson, Daniel Bierwirth, and Patrick Baudisch. 2010. Imaginary interfaces: spatial interaction with empty hands and without visual feedback. In *Proceedings of the 23nd annual ACM symposium on User interface software and technology - UIST '10*, ACM Press, New York, New York, USA, 3–12. DOI:https://doi.org/10.1145/1866029.1866033

[92] Carl Gutwin and Andy Cockburn. 2006. Improving list revisitation with ListMaps. In *Proceedings of the working conference on Advanced Visual Interfaces - AVI '06*, ACM, New York, NY, USA, 396–403. DOI:https://doi.org/10.1145/1133265.1133347

[93] Carl Gutwin, Andy Cockburn, and Nickolas Gough. 2017. A Field Experiment of Spatially-Stable Overviews for Document Navigation. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems - CHI '17*, ACM Press, New York, New York, USA, 5905–5916. DOI:https://doi.org/10.1145/3025453.3025905

[94] Carl Gutwin, Andy Cockburn, and Benjamin Lafreniere. 2015. Testing the rehearsal hypothesis with two FastTap interfaces. In *Proceedings of the 41st Graphics Interface Conference - GI '15*, Canadian Information Processing Society, 223–231. Retrieved May 23, 2018 from https://dl.acm.org/citation.cfm?id=2788930

[95] Carl Gutwin, Andy Cockburn, Joey Scarr, Sylvain Malacria, and Scott C. Olson. 2014. Faster command selection on tablets with FastTap. In *Proceedings of the ACM conference on Human factors in computing systems - CHI '14*, ACM Press, New York, New York, USA, 2617–2626. DOI:https://doi.org/10.1145/2556288.2557136

[96] Sandra G. Hart and Lowell E. Staveland. 1988. Development of NASA-TLX (Task Load Index): Results of Empirical and Theoretical Research. *Advances in Psychology* 52, (1988), 139–183. DOI:https://doi.org/10.1016/S0166-4115(08)62386-9

[97] Lynn Hasher and Rose T. Zacks. 1979. Automatic and effortful processes in memory. *Journal of Experimental Psychology: General* 108, 3 (1979), 356–388. DOI:https://doi.org/10.1037/0096-3445.108.3.356

[98] Mary Hegarty, Daniel R Montello, Anthony E Richardson, Toru Ishikawa, and Kristin Lovelace. 2006. Spatial abilities at different scales: Individual differences in aptitude-test performance and spatial-layout learning. *Intelligence* 34, 2 (March 2006), 151–176. DOI:https://doi.org/10.1016/j.intell.2005.09.005

[99] William E. Hick. 1952. On the rate of gain of information. *Quarterly Journal of Experimental Psychology* 4, 1 (March 1952), 11–26. DOI:https://doi.org/10.1080/17470215208416600

[100] William C. Hill, James D. Hollan, Dave Wroblewski, and Tim McCandless. 1992. Edit wear and read wear. In *Proceedings of the SIGCHI conference on Human factors in computing systems - CHI '92*, ACM Press, New York, NY, USA, 3–9. DOI:https://doi.org/10.1145/142750.142751

[101] Ken Hinckley, Randy Pausch, and Dennis Proffitt. 1997. Attention and visual feedback: the bimanual frame of reference. In *Proceedings of the 1997 symposium on Interactive 3D graphics - SI3D '97*, ACM Press, New York, New York, USA, New York, USA, 121–126. DOI:https://doi.org/10.1145/253284.253318

[102] Ken Hinckley, Randy Pausch, Dennis Proffitt, James Patten, and Neal Kassell. 1997. Cooperative bimanual action. In *Proceedings of the SIGCHI conference on Human factors in computing systems - CHI '97*, ACM Press, New York, New York, USA, 27–34. DOI:https://doi.org/10.1145/258549.258571

[103] Stephen Hirtle. 2009. Cognitive Maps. In *Handbook of Research on Geoinformatics*. IGI Global, 58–64. DOI:https://doi.org/10.4018/978-1-59140-995-3.ch008

[104] Harry Hochheiser and Ben Shneiderman. 2000. Performance benefits of simultaneous over sequential menus as task complexity increases. *International Journal of Human-Computer Interaction* 12, 2 (2000), 173–192. DOI:https://doi.org/10.1207/S15327590IJHC1202_2

[105] Rebecca Hurlebaus, Kai Basten, Hanspeter A. Mallot, and Jan M. Wiener. 2008. Route learning strategies in a virtual cluttered environment. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, Springer Berlin Heidelberg, Berlin, Heidelberg, 104–120. DOI:https://doi.org/10.1007/978-3-540-87601-4_10

[106] Alexandra Ion, Yu Ling Betty Chang, Michael Haller, Mark Hancock, and Stacey D. Scott. 2013. Canyon: providing location awareness of multiple moving objects in a detail view on large displays. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems - CHI '13*, ACM Press, New York, New York, USA, 3149–3158. DOI:https://doi.org/10.1145/2470654.2466431

[107] Toru Ishikawa and Daniel R. Montello. 2006. Spatial knowledge acquisition from direct experience in the environment: Individual differences in the development of metric knowledge and the integration of separately learned places. *Cognitive Psychology* 52, 2 (2006), 93–129. DOI:https://doi.org/https://doi.org/10.1016/j.cogpsych.2005.08.003

[108] Phillip Isola, Devi Parikh, Antonio Torralba, and Aude Oliva. 2011. Understanding the Intrinsic Memorability of Images. In *Proceedings of the 24th International Conference on Neural Information Processing Systems*, Curran Associates, Inc., Granada, Spain, 2429–2437.

[109] Miriam Ittyerah and Lawrence E. Marks. 2007. Memory for curvature of objects: Haptic touch vs. vision. *British Journal of Psychology* 98, 4 (November 2007), 589–610. DOI:https://doi.org/10.1348/000712606X171531

[110] Dan Jackson, James Nicholson, Gerrit Stoeckigt, Rebecca Wrobel, Anja Thieme, and Patrick Olivier. 2013. Panopticon:a parallel video overview system. In *Proceedings of the 26th annual ACM symposium on User interface software and technology - UIST '13*, ACM Press, New York, New York, USA, 123–130. DOI:https://doi.org/10.1145/2501988.2502038

[111] Lucia F. Jacobs. 2003. Spatial Memory. In *Encyclopedia of the Neurological Sciences*. Elsevier, 86–90. DOI:https://doi.org/10.1016/b0-12-226870-9/01375-7

[112] W.Jake Jacobs, Holly E. Laurance, and Kevin G.F. Thomas. 1997. Place Learning in Virtual Space I: Acquisition, Overshadowing, and Transfer. *Learning and Motivation* 28, 4 (November 1997), 521–541. DOI:https://doi.org/10.1006/lmot.1997.0977

[113] Larry L. Jacoby and Dawn Witherspoon. 1982. Remembering without awareness. *Canadian Journal of Psychology/Revue canadienne de psychologie* 36, 2 (1982), 300–324. DOI:https://doi.org/10.1037/h0080638

[114] Petra Jansen-Osmann and Petra Fuchs. 2006. Wayfinding behavior and spatial knowledge of adults and children in a virtual environments: The role of landmarks. *Experimental Psychology* 53, 3 (2006), 171–181. DOI:https://doi.org/10.1027/1618-3169.53.3.171

[115] Yvonne Jansen, Jonas Schjerlund, and Kasper Hornbæk. 2019. Effects of Locomotion and Visual Overview on Spatial Memory when Interacting with Wall Displays. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems - CHI '19*, ACM Press, New York, New York, USA, New York, USA, 1–12. DOI:https://doi.org/10.1145/3290605.3300521

[116] Nikhita Joshi and Daniel Vogel. 2019. An Evaluation of Touch Input at the Edge of a Table. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems - CHI '19*, ACM Press, New York, New York, USA, 1–12. DOI:https://doi.org/10.1145/3290605.3300476

[117] Joseph M. Juran. 1951. *Quality Control Handbook*. McGraw-Hill, New York.

[118] Amy A. Kalia, Paul R. Schrater, and Gordon E. Legge. 2013. Combining Path Integration and Remembered Landmarks When Navigating without Vision. *PLoS ONE* 8, 9 (September 2013), 8. DOI:https://doi.org/10.1371/journal.pone.0072170

[119] Ruth Kanfer and Phillip Ackerman. 2000. Individual Differences in Work Motivation: Further Explorations of a Trait Framework. *Applied Psychology* 49, 3 (July 2000), 470–482. DOI:https://doi.org/10.1111/1464-0597.00026

[120] Victor Kaptelinin. 1993. Item recognition in menu selection: the effect of practice. In *INTERACT '93 and CHI '93 conference companion on Human factors in computing systems - CHI '93*, Association for Computing Machinery (ACM), New York, New York, USA, 183–184. DOI:https://doi.org/10.1145/259964.260196

[121] Roy P. C. Kessels, L Jaap Kappelle, Edward H.F de Haan, and Albert Postma. 2002. Lateralization of spatial-memory processes: evidence on spatial span, maze learning, and memory for object locations. *Neuropsychologia* 40, 8 (January 2002), 1465–1473. DOI:https://doi.org/10.1016/S0028-3932(01)00199-3

[122] Kimia Kiani, George Cui, Andrea Bunt, Joanna McGrenere, and Parmit K. Chilana. 2019. Beyond "One-Size-Fits-All": Understanding the Diversity in How Software Newcomers Discover and Make Use of Help Resources. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, ACM, New York, NY, USA, 1–14. DOI:https://doi.org/10.1145/3290605.3300570

[123] Jong W. Kim, Frank E. Ritter, and Richard J. Koubek. 2013. An integrated theory for improved skill acquisition and retention in the three stages of learning. *Theoretical Issues in Ergonomics Science 14*, 22–37. DOI:https://doi.org/10.1080/1464536X.2011.573008

[124] Juho Kim, Philip J. Guo, Carrie J. Cai, Shang-Wen (Daniel) Li, Krzysztof Z. Gajos, and Robert C. Miller. 2014. Data-driven interaction techniques for improving navigation of educational videos. In *Proceedings of the ACM symposium on User interface software and technology - UIST '14*, ACM, New York, NY, USA, 563–572. DOI:https://doi.org/10.1145/2642918.2647389

[125] Roberta L. Klatzky, Jack M. Loomis, Reginald G. Golledge, Joseph G. Cicinelli, Sally Doherty, and James W. Pellegrino. 1990. Acquisition of route and survey knowledge in the absence of vision. *Journal of Motor Behavior* 22, 1 (1990), 19–43. DOI:https://doi.org/10.1080/00222895.1990.10735500

[126] Gordon Kurtenbach and William Buxton. 1993. The limits of expert performance using hierarchic marking menus. In *Proceedings of the SIGCHI conference on Human factors in computing systems - CHI '93*, ACM Press, New York, New York, USA, 482–487. DOI:https://doi.org/10.1145/169059.169426

[127] Gordon Kurtenbach and William Buxton. 1994. User learning and performance with marking menus. In *Proceedings of the SIGCHI conference on Human factors in computing systems celebrating interdependence - CHI '94*, ACM Press, New York, New York, USA, 258–264. DOI:https://doi.org/10.1145/191666.191759

[128] Gordon Kurtenbach, George W. Fitzmaurice, Russell N. Owen, and Thomas Baudel. 1999. The Hotbox: efficient access to a large number of menu-items. In *Proceedings of the SIGCHI conference on Human factors in computing systems the CHI is the limit - CHI '99*, ACM Press, New York, New York, USA, 231–237. DOI:https://doi.org/10.1145/302979.303047

[129] Gordon Paul Kurtenbach. 1993. The design and evaluation of marking menus. University of Toronto.

[130] Gordon Kurtenbach, Abigail Sellen, and William Buxton. 1993. An Empirical Evaluation of Some Articulatory and Cognitive Aspects of Marking Menus. *Human-Computer Interaction* 8, 1 (March 1993), 1–23. DOI:https://doi.org/10.1207/s15327051hci0801_1

[131] Sari A. Laakso, Karri P. Laakso, and Asko J. Saura. 2000. Improved scroll bars. In *CHI '00 extended abstracts on Human factors in computing systems - CHI '00*, ACM Press, New York, NY, USA, 97–98. DOI:https://doi.org/10.1145/633292.633350

[132] Benjamin Lafreniere, Carl Gutwin, Andy Cockburn, and Tovi Grossman. 2016. Faster Command Selection on Touchscreen Watches. In *Proceedings of the ACM Conference on Human Factors in Computing Systems - CHI '16*, ACM Press, New York, New York, USA, 4663–4674. DOI:https://doi.org/10.1145/2858036.2858166

[133] G.Julian Lepinski, Tovi Grossman, and George Fitzmaurice. 2010. The design and evaluation of multitouch marking menus. In *Proceedings of the ACM conference on Human factors in computing systems - CHI '10*, ACM Press, New York, New York, USA, 2233–2242. DOI:https://doi.org/10.1145/1753326.1753663

[134] Stephen C. Levinson. 1999. Frames of Reference and Molyneux's Question: Crosslinguistic Evidence. In *Language and Space*. The MIT Press. DOI:https://doi.org/10.7551/mitpress/4107.003.0006

[135] John P. Lewis, Ruth Rosenholtz, Nickson Fong, and Ulrich Neumann. 2004. VisualIDs: automatic distinctive icons for desktop interfaces. *ACM Transactions on Graphics* 23, 3 (August 2004), 416–423. DOI:https://doi.org/10.1145/1015706.1015739

[136] Chin-Teng Teng Lin, Teng-Yi Yi Huang, Wen-Jing Jing Lin, Shu-Yen Yen Chang, Yin-Hung Hung Lin, Li-Wei Wei Ko, Daisy L. Hung, and Erik C. Chang. 2012. Gender differences in wayfinding in virtual environments with global or local landmarks. *Journal of Environmental Psychology* 32, 2 (June 2012), 89–96. DOI:https://doi.org/10.1016/j.jenvp.2011.12.004

[137] Chin Teng Lin, Teng Yi Huang, Wen Jing Lin, Shu Yen Chang, Yin Hung Lin, Li Wei Ko, Daisy L. Hung, and Erik C. Chang. 2012. Gender differences in wayfinding in virtual environments with global or local landmarks. *Journal of Environmental Psychology* 32, 2 (June 2012), 89–96. DOI:https://doi.org/10.1016/j.jenvp.2011.12.004

[138] Peter H. Lindsay and Donald A. Norman. 1977. *Human Information Processing*. Elsevier. DOI:https://doi.org/10.1016/C2013-0-11094-X

[139] Robert Lloyd. 1989. Cognitive Maps: Encoding and Decoding Information. *Annals of the Association of American Geographers* 79, 1 (November 1989), 101–124. Retrieved from http://www.jstor.org/stable/2563857

[140] Jack M. Loomis, Roberta L. Klatzky, Reginald G. Golledge, Joseph G. Cicinelli, and et al. 1993. Nonvisual navigation by blind and sighted: Assessment of path integration ability. *Journal of Experimental Psychology: General* 122, 1 (1993), 73–91. DOI:https://doi.org/10.1037//0096-3445.122.1.73

[141] Jack M. Loomis, José A. da Silva, Naofumi Fujita, and Sergio S. Fukusima. 1992. Visual space perception and visually directed action. *Journal of Experimental Psychology: Human Perception and Performance* 18, 4 (1992), 906–921. DOI:https://doi.org/10.1037//0096-1523.18.4.906

[142] Steven Luck. 2007. Visual short term memory. *Scholarpedia* 2, 6 (2007), 3328. DOI:https://doi.org/10.4249/scholarpedia.3328

[143] Kevin Lynch. 1960. *The image of the city*. MIT Press.

[144] Scott MacKenzie. 2013. *Human-Computer Interaction: An Empirical Research Perspective* (1st ed.). Morgan Kaufmann Publishers Inc., San Francisco, CA, USA. Retrieved January 16, 2017 from http://dl.acm.org/citation.cfm?id=2501707

[145] Eleanor A. Maguire, Neil Burgess, and John O'Keefe. 1999. Human spatial navigation: Cognitive maps, sexual dimorphism, and neural substrates. *Current Opinion in Neurobiology* 9, 2 (April 1999), 171–177. DOI:https://doi.org/10.1016/S0959-4388(99)80023-3

[146] Shareen Mahmud, Jessalyn Alvina, Parmit K. Chilana, Andrea Bunt, and Joanna McGrenere. 2020. Learning Through Exploration: How Children, Adults, and Older Adults Interact with a New Feature-Rich Application. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, ACM, New York, NY, USA, 1–14. DOI:https://doi.org/10.1145/3313831.3376414

[147] Aristides Mairena, Carl Gutwin, and Andy Cockburn. 2019. Peripheral Notifications in Large Displays: Effects of Feature Combination and Task Interference. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems - CHI '19*, ACM Press, New York, New York, USA, 1–12. DOI:https://doi.org/10.1145/3290605.3300870

[148] Sylvain Malacria, Gilles Bailly, Joel Harrison, Andy Cockburn, and Carl Gutwin. 2013. Promoting Hotkey use through rehearsal with ExposeHK. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems - CHI '13*, ACM Press, New York, New York, USA, 573–582. DOI:https://doi.org/10.1145/2470654.2470735

[149] Jean M. Mandler, Dale Seegmiller, and Jeanne Day. 1977. On the coding of spatial information. *Memory & Cognition* 5, 1 (January 1977), 10–16. DOI:https://doi.org/10.3758/BF03209185

[150] Justin Matejka, Tovi Grossman, and George Fitzmaurice. 2012. Swift: reducing the effects of latency in online video scrubbing. In *Proceedings of the 2012 ACM annual conference on Human Factors in Computing Systems - CHI '12*, ACM Press, New York, NY, USA, 637–646. DOI:https://doi.org/10.1145/2207676.2207766

[151] Justin Matejka, Tovi Grossman, and George Fitzmaurice. 2013. Swifter: improved online video scrubbing. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems - CHI '13*, ACM Press, New York, NY, USA, 1159–1168. DOI:https://doi.org/10.1145/2470654.2466149

[152] D.Scott McCrickard and Richard Catrambone. 1999. Beyond the scrollbar: an evolution and evaluation of alternative navigation techniques. In *Proceedings 1999 IEEE Symposium on Visual Languages*, IEEE, 270–277. DOI:https://doi.org/10.1109/VL.1999.795913

[153] Timothy P. McNamara, Björn Rump, and Steffen Werner. 2003. Egocentric and geocentric frames of reference in memory of large-scale space. *Psychonomic Bulletin and Review* 10, 3 (2003), 589–595. DOI:https://doi.org/10.3758/BF03196519

[154] Ehsan Sotoodeh Mollashahi, Md. Sami Uddin, and Carl Gutwin. 2018. Improving revisitation in long documents with two-level artificial-landmark scrollbars. In *Proceedings of the 2018 International Conference on Advanced Visual Interfaces*, ACM, New York, NY, USA, 1–9. DOI:https://doi.org/10.1145/3206505.3206554

[155] Daniel R. Montello. 1998. A New Framework for Understanding the Acquisition of Spatial Knowledge in Large-Scale Environments. In *Spatial and temporal reasoning in geographic information systems*, M. J. Egenhofer and R. G. Golledge (eds.). Oxford University Press., New York, 143–154.

[156] Weimin Mou and Timothy P. McNamara. 2002. Intrinsic Frames of Reference in Spatial Memory. *Journal of Experimental Psychology: Learning Memory and Cognition* 28, 1 (2002), 162–170. DOI:https://doi.org/10.1037/0278-7393.28.1.162

[157] Weimin Mou, Chengli Xiao, and Timothy P. McNamara. 2008. Reference directions and reference objects in spatial memory of a briefly viewed layout. *Cognition* 108, 1 (July 2008), 136–154. DOI:https://doi.org/10.1016/j.cognition.2008.02.004

[158] Bennet B. Murdock. 1962. The serial position effect of free recall. *Journal of Experimental*

Psychology 64, 5 (November 1962), 482–488. DOI:https://doi.org/10.1037/h0045106

[159] Allen Newell and Paul S. Rosenbloom. 1981. Mechanisms of skill acquisition and the law of practice. *Cognitive skills and their acquisition* 1, (1981), 1–55.

[160] Annie Wy Ng and Alan Hs Chan. 2009. What makes an icon effective? In *AIP Conference Proceedings*, AIP, 104–114. DOI:https://doi.org/10.1063/1.3078113

[161] Jakob Nielsen. 1993. *Usability Engineering*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.

[162] Jakob Nielsen. 1994. Enhancing the explanatory power of usability heuristics. In *Proceedings of the SIGCHI conference on Human factors in computing systems celebrating interdependence - CHI '94* (CHI '94), ACM Press, New York, New York, USA, 152–158. DOI:https://doi.org/10.1145/191666.191729

[163] Raffaella Nori and Laura Piccardi. 2011. Familiarity and spatial cognitive style: How important are they for spatial representation? In *Spatial Memory: Visuospatial Processes, Cognitive Performance and Developmental Effects*. Nova Science Publishers, Inc., 123–144.

[164] Kenton O'Hara and Abigail Sellen. 1997. Comparison of reading paper and on-line documents. In *Conference on Human Factors in Computing Systems - Proceedings*, ACM, New York, New York, USA, 335–342. DOI:https://doi.org/10.1145/258549.258787

[165] Daniel L. Odell, Richard C. Davis, Andrew Smith, and Paul K. Wright. 2004. Toolglasses, marking menus, and hotkeys: a comparison of one and two-handed command selection techniques. In *Proceedings of Graphics Interface 2004 - GI '04*, CHCCS, Waterloo, Ontario, Canada, 17–24. Retrieved September 8, 2016 from http://dl.acm.org/citation.cfm?id=1006058.1006061

[166] John R. Pani and Davido Dupree. 1994. Spatial Reference Systems in the Comprehension of Rotational Motion. *Perception* 23, 8 (August 1994), 929–946. DOI:https://doi.org/10.1068/p230929

[167] Amy Pavel, Dan B. Goldman, Björn Hartmann, and Maneesh Agrawala. 2015. SceneSkim: Searching and Browsing Movies Using Synchronized Captions, Scripts and Plot Summaries. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software & Technology - UIST '15*, ACM Press, New York, NY, USA, 181–190. DOI:https://doi.org/10.1145/2807442.2807502

[168] Amy Pavel, Colorado Reed, Björn Hartmann, and Maneesh Agrawala. 2014. Video Digests: a browsable, skimmable format for informational lecture videos. In *Proceedings of the 27th annual ACM symposium on User interface software and technology - UIST '14*, ACM Press, New York, NY, USA, 573–582. DOI:https://doi.org/10.1145/2642918.2647400

[169] Simon T. Perrault, Eric Lecolinet, Yoann Pascal Bourse, Shengdong Zhao, and Yves Guiard. 2015. Physical Loci: leveraging spatial, object and semantic memory for command selection. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems - CHI '15*, ACM, New York, NY, USA, 299–308. DOI:https://doi.org/10.1145/2702123.2702126

[170] Albert Postma and Edward H.F. De Haan. 1996. What was where? Memory for object locations. *The Quarterly journal of experimental psychology. A, Human experimental psychology* 49, 1 (February 1996), 178–99. DOI:https://doi.org/10.1080/713755605

[171] Albert Postma, Roy P. C. Kessels, and Marieke van Asselen. 2004. The Neuropsychology of Object-Location Memory. In *Human spatial memory: Remembering where.* Lawrence Erlbaum Associates Publishers, Postma, Albert: Psychological Laboratory, Helmholtz Instituut, Utrecht University, Heidelberglaan 2, Utrecht, Netherlands, 3584 CS, 143–160.

[172] Albert Postma, Roy P. C. Kessels, and Marieke van Asselen. 2008. How the brain remembers and forgets where things are: the neurocognition of object-location memory. *Neuroscience and biobehavioral reviews* 32, 8 (October 2008), 1339–45. DOI:https://doi.org/10.1016/j.neubiorev.2008.05.001

[173] Clark C. Presson and Daniel R. Montello. 1988. Points of reference in spatial cognition: Stalking the elusive landmark. *British Journal of Developmental Psychology* 6, 4 (November 1988), 378–381. DOI:https://doi.org/10.1111/j.2044-835x.1988.tb01113.x

[174] Teriitutea Quesnot and Stéphane Roche. 2014. Measure of Landmark Semantic Salience through Geosocial Data Streams. *ISPRS International Journal of Geo-Information* 4, 1 (December 2014), 1–31. DOI:https://doi.org/10.3390/ijgi4010001

[175] Philip Quinn, Andy Cockburn, Indratmo, and Carl Gutwin. 2008. An investigation of dynamic landmarking functions. In *Proceedings of the working conference on Advanced visual interfaces - AVI '08*, ACM Press, New York, New York, USA, New York, USA, 322. DOI:https://doi.org/10.1145/1385569.1385623

[176] John Rieman. 1996. A field study of exploratory learning strategies. *ACM Transactions on Computer-Human Interaction* 3, 3 (September 1996), 189–218. DOI:https://doi.org/10.1145/234526.234527

[177] George Robertson, Mary Czerwinski, Kevin Larson, Daniel C. Robbins, David Thiel, and Maarten van Dantzich. 1998. Data mountain: using spatial memory for document management. In *Proceedings of the ACM symposium on User interface software and technology - UIST '98*, ACM Press, New York, New York, USA, 153–162. DOI:https://doi.org/10.1145/288392.288596

[178] Irvin Rock and Phima Engelstein. 1959. A Study of Memory for Visual Form. *The American Journal of Psychology* 72, 2 (June 1959), 229. DOI:https://doi.org/10.2307/1419366

[179] Beverly Roskos-Ewoldsen, Timothy P. McNamara, Amy L. Shelton, and Walter Carr. 1998. Mental representations of large and small spatial layouts are orientation dependent. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 24, 1 (1998), 215–226. DOI:https://doi.org/10.1037//0278-7393.24.1.215

[180] Roy A. Ruddle, Ekaterina Volkova, and Heinrich H. Bülthoff. 2013. Learning to walk in virtual reality. *ACM Transactions on Applied Perception* 10, 2 (May 2013), 1–17. DOI:https://doi.org/10.1145/2465780.2465785

[181] Edward K. Sadalla, W. Jeffrey Burroughs, and Lorin J. Staplin. 1980. Reference points in spatial cognition. *Journal of Experimental Psychology: Human Learning and Memory* 6, 5 (September 1980), 516–528. DOI:https://doi.org/10.1037/0278-7393.6.5.516

[182] Krystian Samp. 2013. Designing graphical menus for novices and experts: connecting design characteristics with design goals. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM, New York, NY, USA, 3159–3168. DOI:https://doi.org/10.1145/2470654.2466432

[183] Joey Scarr, Andy Cockburn, Carl Gutwin, and Andrea Bunt. 2012. Improving command selection with CommandMaps. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems - CHI '12*, ACM, New York, NY, USA, 257–266. DOI:https://doi.org/10.1145/2207676.2207713

[184] Joey Scarr, Andy Cockburn, Carl Gutwin, Andrea Bunt, Jared E. Cechanowicz, Joey Scarr, Andy Cockburn, Carl Gutwin, Andrea Bunt, and Jared E. Cechanowicz. 2014. The usability of CommandMaps in realistic tasks. In *Proceedings of the 32nd annual ACM conference on Human factors in computing systems - CHI '14*, ACM Press, New York, New York, USA, 2241–2250. DOI:https://doi.org/10.1145/2556288.2556976

[185] Joey Scarr, Andy Cockburn, Carl Gutwin, and Sylvain Malacria. 2013. Testing the robustness and performance of spatially consistent interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM, 3139–3148. DOI:https://doi.org/10.1145/2470654.2466430

[186] Joey Scarr, Andy Cockburn, Carl Gutwin, and Philip Quinn. 2011. Dips and ceilings: understanding and supporting transitions to expertise in user interfaces. In *Proceedings of the 2011 annual conference on Human factors in computing systems - CHI '11*, ACM Press, New York, New York, USA, 2741–2750. DOI:https://doi.org/10.1145/1978942.1979348

[187] Joseph Laurence Scarr. 2014. Understanding and Exploiting Spatial Memory in the Design of Efficient Command Selection Interfaces. University of Canterbury. Retrieved August 25, 2016 from http://ir.canterbury.ac.nz/handle/10092/9326

[188] Richard A. Schmidt and Timothy Donald Lee. 2011. *Motor control and learning: A behavioral emphasis* (5th ed.). Human Kinetics, Champaign, IL, US.

[189] Walter Schneider and Richard M. Shiffrin. 1977. Controlled and automatic human information processing: I. Detection, search, and attention. *Psychological Review* 84, 1 (January 1977), 1–66. DOI:https://doi.org/10.1037/0033-295X.84.1.1

[190] Klaus Schoeffmann, Mario Taschwer, and Laszlo Boeszoermenyi. 2010. The video explorer: a tool for navigation and searching within a single video based on fast content analysis. In *Proceedings of the ACM SIGMM conference on Multimedia systems - MMSys '10*, ACM, New York, NY, USA, 247. DOI:https://doi.org/10.1145/1730836.1730867

[191] Katherine Schramm, Carl Gutwin, and Andy Cockburn. 2016. Supporting Transitions to Expertise in Hidden Toolbars. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems - CHI '16*, ACM, New York, NY, USA, 4687–4698. DOI:https://doi.org/10.1145/2858036.2858412

[192] Amy L. Shelton and Timothy P. McNamara. 1997. Multiple views of spatial memory. *Psychonomic Bulletin and Review* 4, 1 (1997), 102–106. DOI:https://doi.org/10.3758/BF03210780

[193] Amy L. Shelton and Timothy P. McNamara. 2001. Systems of Spatial Reference in Human Memory. *Cognitive Psychology* 43, 4 (December 2001), 274–310. DOI:https://doi.org/10.1006/cogp.2001.0758

[194] Ben Shneiderman, Catherine Plaisant, Maxine Cohen, Steven Jacobs, Niklas Elmqvist, and Nicholas Diakopoulos. 2018. *Designing the User Interface: Strategies for Effective Human-Computer Interaction* (6th ed.). Pearson. Retrieved April 21, 2020 from /content/one-dot-com/one-dot-com/us/en/higher-education/program.html

[195] Alexander W. Siegel and Sheldon H. White. 1975. The Development of Spatial Representations of Large-Scale Environments. *Advances in Child Development and Behavior* 10, (January 1975), 9–55. DOI:https://doi.org/10.1016/S0065-2407(08)60007-5

[196] Robert S. Siegler. 1991. *Children's thinking*. Prentice-Hall, Inc.

[197] Herbert A. A. Simon. 1987. Satisficing. In *The New Palgrave Dictionary of Economics*, P. Eatwell,J., Milgate, M., and Newman (ed.). Stockton Press, New York, 243–245. DOI:https://doi.org/10.1057/9780230226203.3474

[198] Amy Skopik and Carl Gutwin. 2005. Improving revisitation in fisheye views with visit wear. In *Proceedings of the SIGCHI conference on Human factors in computing systems - CHI '05*, ACM Press, New York, New York, USA, 771–780. DOI:https://doi.org/10.1145/1054972.1055079

[199] George S. Snoddy. 1926. Learning and stability: a psychophysiological analysis of a case of motor learning with clinical applications. *Journal of Applied Psychology* 10, 1 (March 1926), 1–36. DOI:https://doi.org/10.1037/h0075814

[200] Molly E. Sorrows and Stephen C. Hirtle. 1999. The Nature of Landmarks for Real and Electronic Spaces. In *Spatial Information Theory. Cognitive and Computational Foundations of Geographic Information Science*. Springer Berlin Heidelberg, 37–50. DOI:https://doi.org/10.1007/3-540-48384-5_3

[201] George Sperling. 1960. The information available in brief visual presentations. *Psychological Monographs: General and Applied* 74, 11 (1960), 1–29. DOI:https://doi.org/10.1037/h0093759

[202] Lionel Standing. 1973. Learning 10,000 pictures. *The Quarterly journal of experimental psychology* 25, 2 (May 1973), 207–22. DOI:https://doi.org/10.1080/14640747308400340

[203] Sibylle D. Steck and Hanspeter A. Mallot. 2000. The Role of Global and Local Landmarks in Virtual Environment Navigation. *Presence: Teleoperators and Virtual Environments* 9, 1 (February 2000), 69–83. DOI:https://doi.org/10.1162/105474600566628

[204] Jeffery J. Summers. 1989. Motor programs. In *Human skills, 2nd ed.* John Wiley & Sons, Oxford, England, 49–69.

[205] Ke Sun, Chun Yu, and Yuanchun Shi. 2019. Exploring Low-Occlusion Qwerty Soft Keyboard Using Spatial Landmarks. *ACM Transactions on Computer-Human Interaction* 26, 4 (June 2019), 1–33. DOI:https://doi.org/10.1145/3318141

[206] Linda Tauscher and Saul Greenberg. 1997. How people revisit web pages: empirical findings and implications for the design of history systems. *International Journal of Human-Computer Studies* 47, 1 (July 1997), 97–137. DOI:https://doi.org/10.1006/ijhc.1997.0125

[207] Caitlin Tenison and John R. Anderson. 2016. Modeling the distinct phases of skill acquisition. *Journal of Experimental Psychology: Learning Memory and Cognition* 42, 5 (May 2016), 749–767. DOI:https://doi.org/10.1037/xlm0000204

[208] Perry W. Thorndyke and Sarah E. Goldin. 1983. Spatial Learning and Reasoning Skill. In *Spatial Orientation*. Springer US, Boston, MA, MA, 195–217. DOI:https://doi.org/10.1007/978-1-4615-9325-6_9

[209] Perry W. Thorndyke and Barbara Hayes-Roth. 1982. Differences in spatial knowledge acquired from maps and navigation. *Cognitive Psychology* 14, 4 (October 1982), 560–589. DOI:https://doi.org/10.1016/0010-0285(82)90019-6

[210] Michael Tlauka and Paul N. Wilson. 1994. The effect of landmarks on route-learning in a computer-simulated environment. *Journal of Environmental Psychology* 14, 4 (1994), 305–313. DOI:https://doi.org/https://doi.org/10.1016/S0272-4944(05)80221-X

[211] Dejan Todorovic. 2008. Gestalt principles. *Scholarpedia* 3, 12 (2008), 5345. DOI:https://doi.org/10.4249/scholarpedia.5345

[212] Edward C. Tolman. 1948. Cognitive maps in rats and men. *Psychological Review* 55, 4 (1948), 189–208. DOI:https://doi.org/10.1037/h0061626

[213] Phillip D. Tomporowski. 2003. *The Psychology of Skill: A Life-Span Approach*. Praeger.

[214] Endel Tulving. 1972. Episodic and semantic memory. *Organization of memory* 1, (1972), 381–403.

[215] Md. Sami Uddin. 2016. Improving Multi-Touch Interactions Using Hands as Landmarks. University of Saskatchewan. Retrieved from http://hci.usask.ca/publications/view.php?id=401

[216] Md. Sami Uddin. 2016. Use of Landmarks to Design Large and Efficient Command Interfaces. In *Proceedings of the 2016 ACM Companion on Interactive Surfaces and Spaces - ISS Companion '16*, ACM Press, New York, New York, USA, 13–17. DOI:https://doi.org/10.1145/3009939.3009942

[217] Md. Sami Uddin and Carl Gutwin. 2016. Rapid Command Selection on Multi-Touch Tablets with Single-Handed HandMark Menus. In *Proceedings of the 2016 ACM International Conference on Interactive Surfaces and Spaces*, ACM, New York, NY, USA, 205–214. DOI:https://doi.org/10.1145/2992154.2992172

[218] Md. Sami Uddin, Carl Gutwin, and Andy Cockburn. 2017. The Effects of Artificial Landmarks on Learning and Performance in Spatial-Memory Interfaces. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, ACM, New York, NY, USA, 3843–3855. DOI:https://doi.org/10.1145/3025453.3025497

[219] Md. Sami Uddin, Carl Gutwin, and Alix Goguey. 2017. Using artificial landmarks to improve revisitation performance and spatial learning in linear control widgets. In *Proceedings of the 5th Symposium on Spatial User Interaction*, ACM, New York, NY, USA, 48–57. DOI:https://doi.org/10.1145/3131277.3132184

[220] Md. Sami Uddin, Carl Gutwin, and Benjamin Lafreniere. 2016. HandMark Menus: Rapid Command Selection and Large Command Sets on Multi-Touch Displays. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, ACM, New York, NY, USA, 5836–5848. DOI:https://doi.org/10.1145/2858036.2858211

[221] Michael T. Ullman. 2004. Contributions of memory circuits to language: the declarative/procedural model. *Cognition* 92, 1–2 (May 2004), 231–270. DOI:https://doi.org/10.1016/j.cognition.2003.10.008

[222] Norman G. Vinson. 1999. Design guidelines for landmarks to support navigation in virtual environments. In *Proceedings of the SIGCHI conference on Human factors in computing systems the CHI is the limit - CHI '99*, ACM Press, New York, New York, USA, 278–285. DOI:https://doi.org/10.1145/302979.303062

[223] David Waller and Yvonne Lippa. 2007. Landmarks as beacons and associative cues: Their role in route learning. *Memory and Cognition* 35, 5 (July 2007), 910–924. DOI:https://doi.org/10.3758/BF03193465

[224] Steffen Werner, Bernd Krieg-Brückner, Hanspeter A. Mallot, Karin Schweizer, and Christian Freksa. 1997. Spatial Cognition: The Role of Landmark, Route, and Survey Knowledge in Human and Robot Navigation. . Springer, Berlin, Heidelberg, 41–50. DOI:https://doi.org/10.1007/978-3-642-60831-5_8

[225] Steffen Werner and Kristine Schmidt. 1999. Environmental reference systems for large-scale spaces. *Spatial Cognition and Computation* 1, 4 (1999), 447–473. DOI:https://doi.org/10.1023/A:1010095831166

[226] Arthur Wingfield. 1979. *Human Learning and Memory: An Introduction* (1st ed.). Harpercollins College Div.

[227] Johan M. von Wright, P. Gebhard, and M. Karttunen. 1975. A developmental study of the recall of spatial location. *Journal of Experimental Child Psychology* 20, 1 (August 1975), 181–190. DOI:https://doi.org/10.1016/0022-0965(75)90037-5

[228] Haijun Xia, Ken Hinckley, Michel Pahud, Xiao Tu, and Bill Buxton. 2017. WritLarge: Ink Unleashed by Unified Scope, Action, & Zoom. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems - CHI '17*, ACM Press, New York, New York, USA, New York, USA, 3227–3240. DOI:https://doi.org/10.1145/3025453.3025664

[229] Yukang Yan, Chun Yu, Xiaojuan Ma, Shuai Huang, Hasan Iqbal, and Yuanchun Shi. 2018. Eyes-Free Target Acquisition in Interaction Space around the Body for Virtual Reality. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems - CHI '18*, ACM Press, New York, New York, USA, New York, USA, 1–13. DOI:https://doi.org/10.1145/3173574.3173616

[230] Frances A. Yates. 1966. *The art of memory*. University of Chicago Press, Chicago.

[231] Bin Yu, Wei-Ying Ma, Klara Nahrstedt, and Hong-Jiang Zhang. 2003. Video summarization based on user log enhanced link analysis. In *Proceedings of the eleventh ACM international conference on Multimedia - MULTIMEDIA '03*, ACM Press, New York, New York, USA, 382–391. DOI:https://doi.org/10.1145/957013.957095

[232] Polle T. Zellweger, Jock D. Mackinlay, Lance Good, Mark Stefik, and Patrick Baudisch. 2003. City lights: contextual views in minimal space. In *CHI '03 extended abstracts on Human factors in computing systems - CHI '03*, ACM Press, New York, New York, USA, 838–839. DOI:https://doi.org/10.1145/765891.766022

[233] Shumin Zhai and Per-Ola Kristensson. 2003. Shorthand writing on stylus keyboard. In *Proceedings of the conference on Human factors in computing systems - CHI '03*, ACM Press, New York, New York, USA, 97–104. DOI:https://doi.org/10.1145/642611.642630

[234] Jingjie Zheng, Xiaojun Bi, Kun Li, Yang Li, and Shumin Zhai. 2018. M3 Gesture Menu: Design and Experimental Analyses of Marking Menus for Touchscreen Mobile Interaction. In *Proceedings of the Conference on Human Factors in Computing Systems - CHI '18*, ACM Press, New York, NY, USA, 1–14. DOI:https://doi.org/10.1145/3173574.3173823

[235] Jingjie Zheng, Blaine Lewis, Jeff Avery, and Daniel Vogel. 2018. FingerArc and FingerChord: Supporting Novice to Expert Transitions with Guided Finger-Aware Shortcuts. In *The 31st Annual ACM Symposium on User Interface Software and Technology - UIST '18*, ACM Press, New York, New York, USA, 347–363. DOI:https://doi.org/10.1145/3242587.3242589

[236] George Kingsley Zipf. 1949. *Human Behavior and the Principle of Least Effort: An Introduction to Human Ecology* (Reading, M ed.). Addison-Wesley.

[237] Microsoft Word. Retrieved August 10, 2018 from https://www.products.office.com/word

[238] Facebook. Retrieved August 10, 2018 from https://www.facebook.com/

[239] Adobe Acrobat Reader. Retrieved August 10, 2018 from https://acrobat.adobe.com/us/en/acrobat/pdf-reader.html

[240] Adobe Photoshop. Retrieved August 10, 2018 from https://www.adobe.com/products/photoshop.html

[241] What are Gestalt Principles? | Interaction Design Foundation. Retrieved September 6, 2020 from https://www.interaction-design.org/literature/topics/gestalt-principles

[242] Facebook's old web design will disappear in September - The Verge. Retrieved September 12, 2020 from https://www.theverge.com/2020/8/21/21395079/facebook-new-design-default-september-classic-interface-disappearing

[243] New Facebook Design not showing up for me, is there anything I can do? : facebook. Retrieved September 12, 2020 from https://www.reddit.com/r/facebook/comments/g07zmv/new_facebook_design_not_showing_up_for_me_is/

[244] 10 Heuristics for User Interface Design: Article by Jakob Nielsen. Retrieved September 6, 2020 from https://www.nngroup.com/articles/ten-usability-heuristics/

[245] Method of loci. Retrieved February 20, 2018 from http://en.wikipedia.org/wiki/Method_of_loci

[246] Downloads | Android Developers. Retrieved September 19, 2016 from http://developer.android.com/design/downloads/index.html

[247] Thesis Film - Last Shot: A. Widodo. Retrieved May 1, 2017 from

http://awidodoportfolio.blogspot.ca/p/film.html

[248] 2017. WHO Global report on diabetes. Retrieved May 1, 2017 from http://www.who.int/diabetes/global-report/en/

APPENDIX A FOR MANUSCRIPT A

A.1 STUDY CONSENT FORMS

UNIVERSITY OF SASKATCHEWAN

Department of Computer Science
176 Thorvaldson Building
110 Science Place Saskatoon SK S7N 5C9 Canada
Telephone: (306) 966-4886 Facsimile: (306) 966-4884

Participant Consent Form

You are invited to participate in a research study entitled: The Image of the Interface

Researcher(s): Md. Sami Uddin, PhD Student, Computer Science, sami.uddin@usask.ca, (306) 966-2827

Supervisor: Dr. Carl Gutwin, Professor, Department of Computer Science, gutwin@cs.usask.ca, (306) 966-8646

Purpose(s) and Objective(s) of the Research:

- This study is concerned with understanding the memorability of the visual forms of the computer interfaces that we use. We are interested to see how people use their spatial memory, particularly the available landmarks, to remember different commands in computer interfaces.

Procedures:

- The session will require about 60 minutes, during which you will be asked to answer questions regarding the locations of different commands in an interface from your experience. You will be asked to draw a simple sketch of that interface and to navigate through that interface in the Human-Computer Interaction Lab at the University of Saskatchewan.
- The session will consist of filling out questionnaires, a verbal interview, and performing simple tasks on the sketch of the interface, which will be observed by the experimenter.
- You will be audio-recorded during the interview. No third party other than the researchers involved in this study will have access to the audio recordings. You can have the audio recording device turned off at any time without giving a reason.
- You will also be asked to provide a screenshot of an interface from your frequently used device. If we wish to use your screenshot in publications, we will recreate the configuration and remove all identifiable information.
- Please feel free to ask any questions regarding the procedures and goals of the study or your role.

Funded by: Natural Sciences and Engineering Research Council of Canada (NSERC).

Potential Risks:

- There are no known or anticipated risks to you by participating in this research.
- At the end of the session, you will be given more information about the purpose and goals of the study, and there will be time for you to ask questions about the research.

Potential Benefits:

- Your experience with interfaces from a spatial memory perspective will assist designers in improving the memorability and usability of graphical user interfaces.

Compensation:

- You will receive a $10 honorarium at the end of the session as one way of thanking you for your time.

Confidentiality:

- All data will be kept confidential. Confidentiality will be preserved by using pseudonyms in any presentation of textual data in journals or at conferences.
- **Storage of Data:** The informed consent form and all research data will be kept in a secure location under confidentiality in accordance with University policy for 5 years post publication. After that data will be destroyed completely.

Right to Withdraw:

- Your participation is voluntary, and you can answer only those questions that you are comfortable with. You may withdraw from the research project for any reason, at any time, without explanation; there is no penalty or loss of any reward if you withdraw.
- Should you wish to withdraw, your data will be deleted from the study and destroyed. Your right to withdraw data from the study will apply until August 31, 2018. After this, it is possible that some form of research dissemination will have already occurred, and it may not be possible to withdraw your data.

Follow up:

- To obtain results from the study, please email sami.uddin@usask.ca two months after the study has been completed. The results will also be available on the HCI lab's website: http://www.hci.usask.ca/

Questions or Concerns:

- Contact the researcher(s) using the information at the top of page 1
- This research project has been approved on ethical grounds by the University of Saskatchewan Research Ethics Board. Any questions regarding your rights as a participant may be addressed to that committee through the Research Ethics Office ethics.office@usask.ca (306) 966-2975. Out of town participants may call toll free (888) 966-2975.

Consent:

Your signature below indicates that you have read and understand the description provided.

I have had an opportunity to ask questions and my questions have been answered. I consent to participate in the research project. A copy of this Consent Form has been given to me for my records.

____________________ ____________________ ____________

Name of Participant *Signature* *Date*

________________ ____________

Researcher's Signature *Date*

A copy of this consent will be left with you, and a copy will be taken by the researcher.

A.2 QUESTIONNAIRE

Demographics and interface usage questionnaire.

1. Participant ID: ________
2. Age (in years): ________
3. Gender
   a. Male
   b. Female
   c. Other

4. What is your profession? ______________________________

5. How familiar are you with the interface of __________ (Microsoft Word/Adobe Acrobat Reader/ Photoshop/Facebook; note that only one was used for each participant)?
   - Unfamiliar
   - Not so familiar
   - Somewhat familiar
   - Very familiar
   - Extremely familiar

6. How long have you been using the application (in years)?
   - Less than 1
   - 1-3
   - 4-6
   - 7-9
   - 10+

7. How frequently do you use the application?
   - Everyday
   - A few times a week
   - About once a week
   - A few times a month
   - Once a month
   - Less than once a month

8. On what device do you mostly use the application (e.g., desktop, laptop, tablet, or smartphone)? If you use multiple devices, please name them all.

9. Do you use shortcuts for any command?
   a. Yes. How many (approximately)? _________
   b. No.

10. Please name other application interfaces that you frequently use, and state the device on which you typically use those applications.

A.3 WASHEDOUT IMAGES

Microsoft Word Interface (Home and Insert tabs):

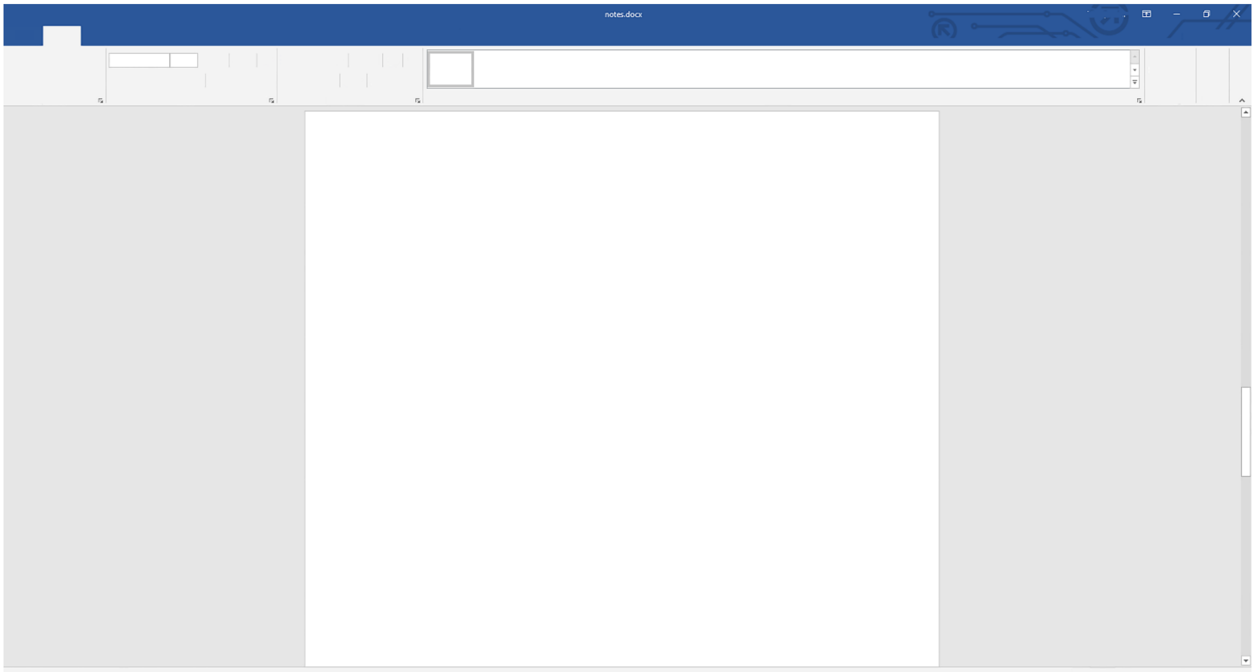

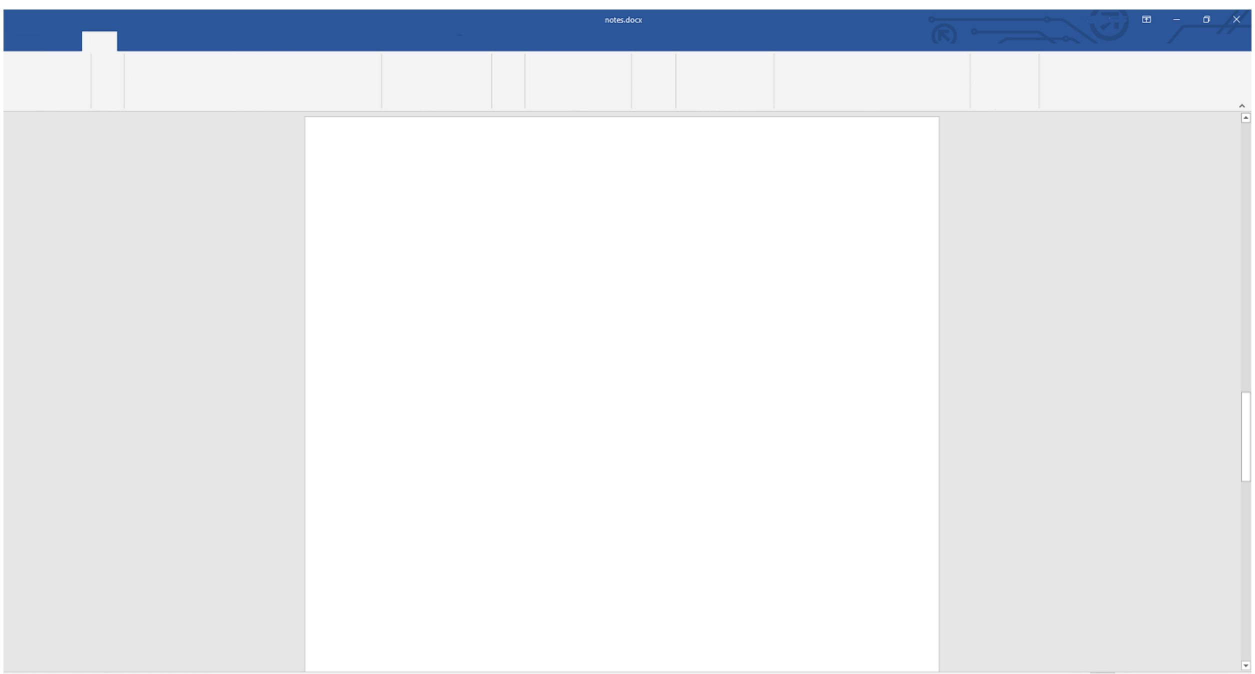

Facebook Interface (Homepage and Profile):

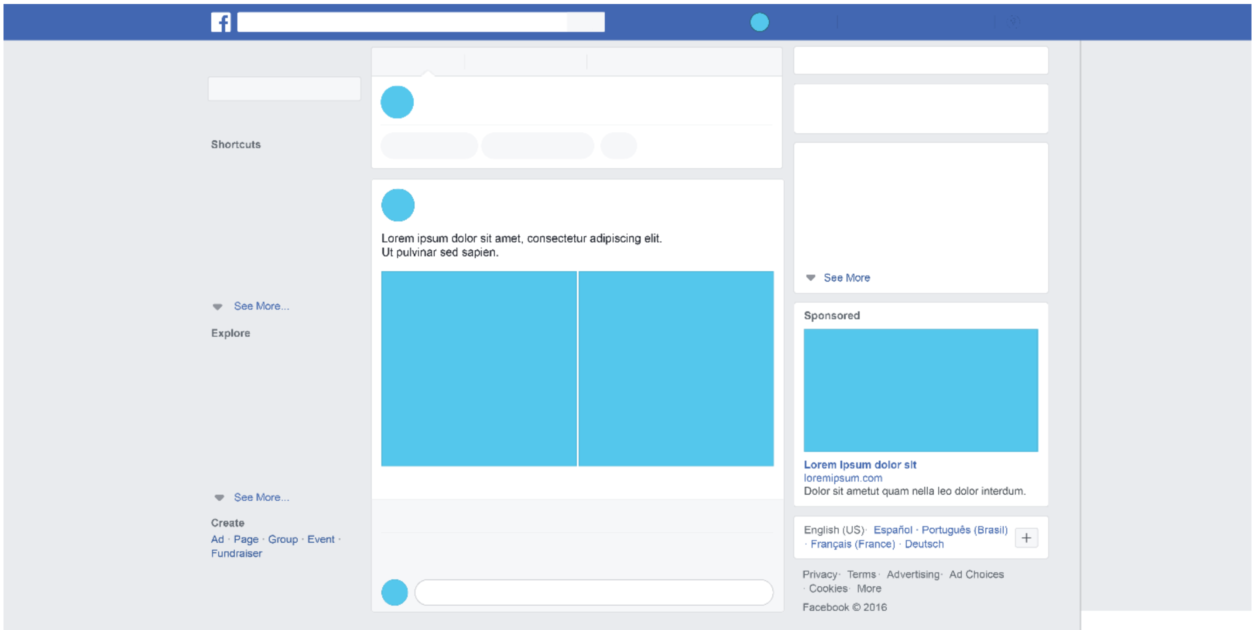

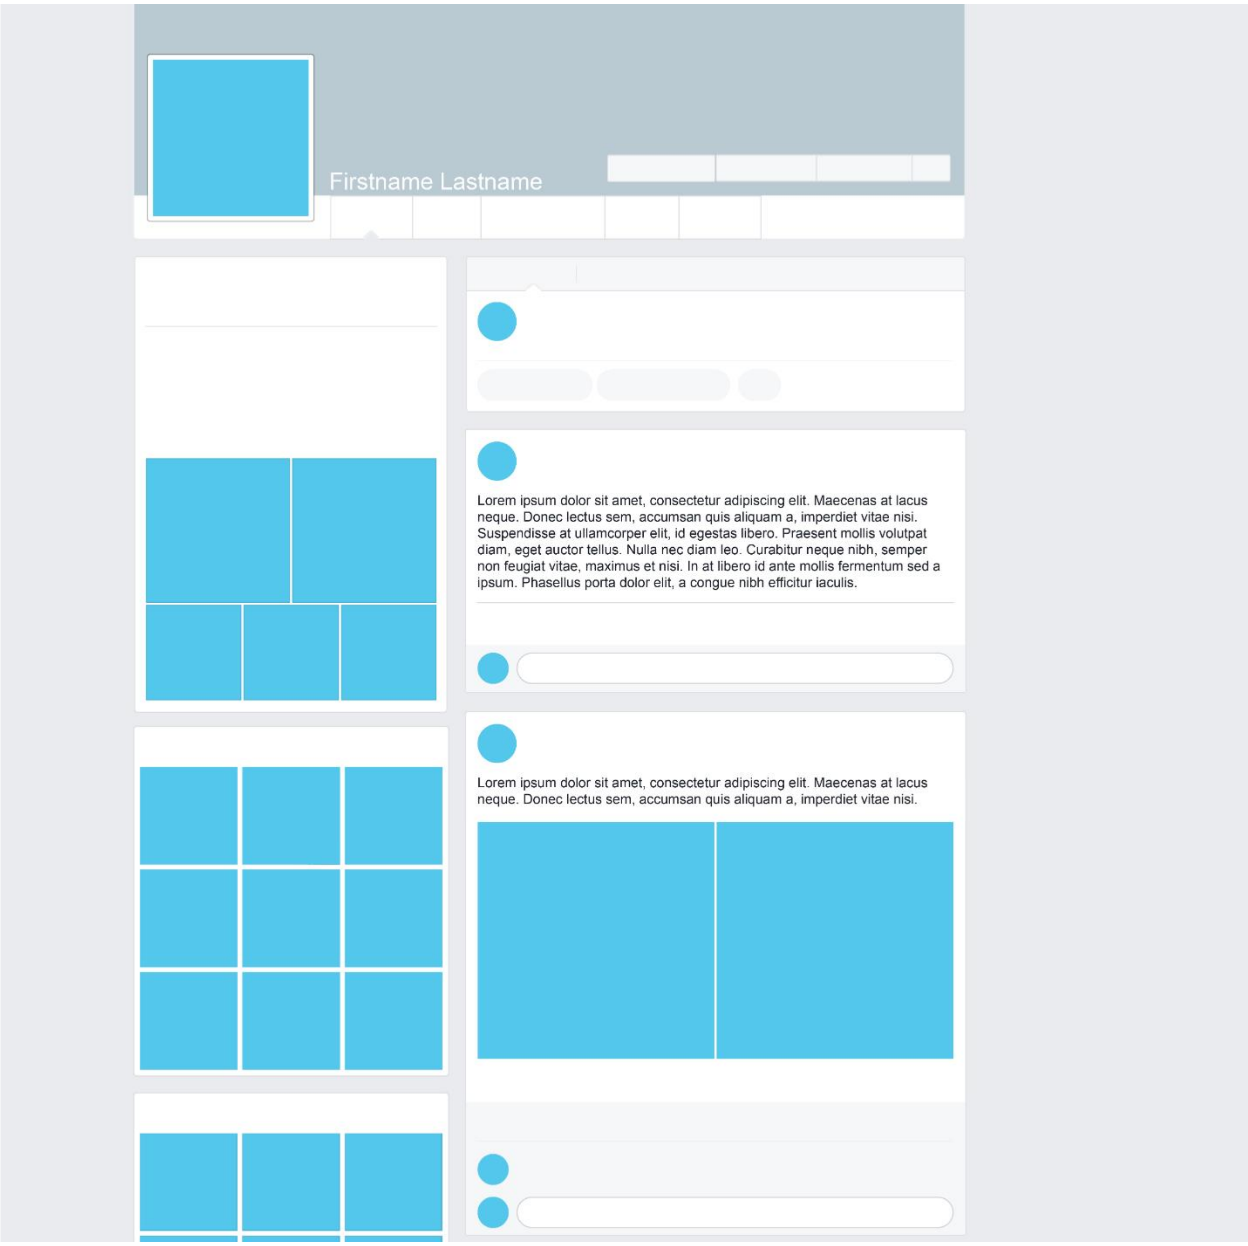
Firstname Lastname
Lorem ipsum dolor sit amet, consectetur adipiscing elit. Maecenas at lacus neque. Donec lectus sem, accumsan quis aliquam a, imperdiet vitae nisi. Suspendisse at ullamcorper elit, id egestas libero. Praesent mollis volutpat diam, eget auctor tellus. Nulla nec diam leo. Curabitur neque nibh, semper non feugiat vitae, maximus et nisi. In at libero id ante mollis fermentum sed a ipsum. Phasellus porta dolor elit, a congue nibh efficitur iaculis.
Lorem ipsum dolor sit amet, consectetur adipiscing elit. Maecenas at lacus neque. Donec lectus sem, accumsan quis aliquam a, imperdiet vitae nisi.

Acrobat Reader Interface:

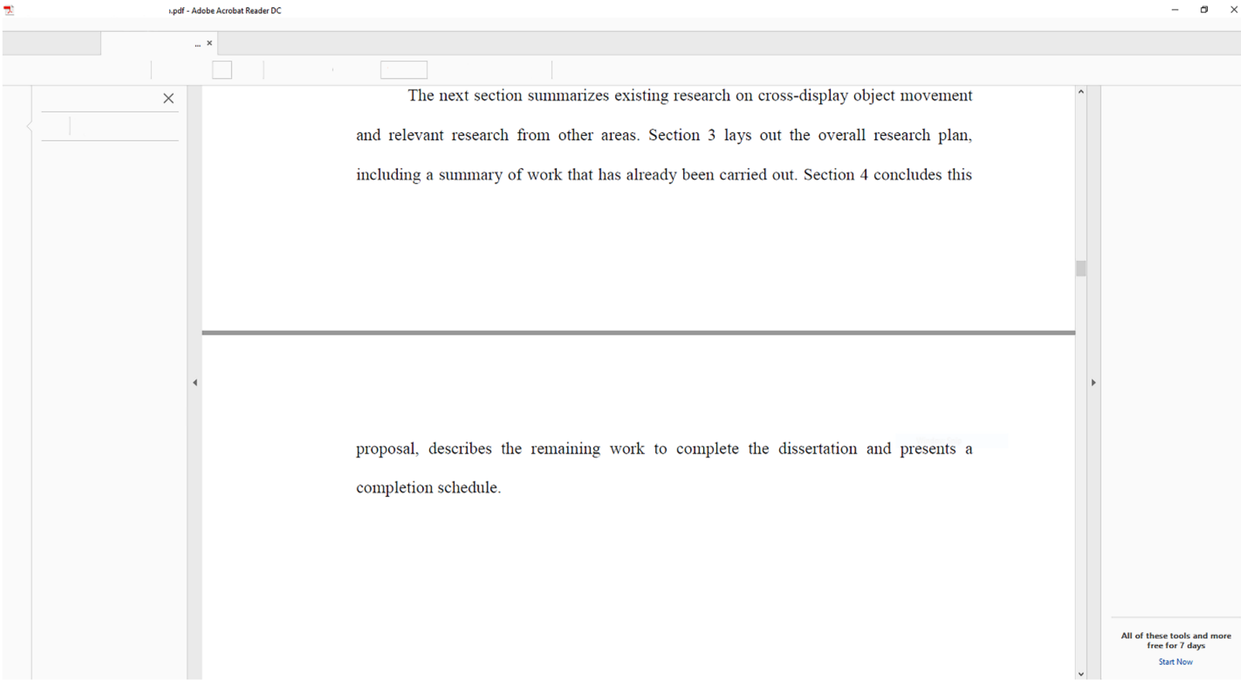

Photoshop Interfaces:

APPENDIX B FOR MANUSCRIPT B

B.1 STUDY CONSENT FORMS

Consent form used for small interfaces.

DEPARTMENT OF COMPUTER SCIENCE
UNIVERSITY OF SASKATCHEWAN
INFORMED CONSENT FORM

UNIVERSITY OF SASKATCHEWAN

Research Project: **Artificial Landmark Small Study – Spring 2016**

Investigators: Dr. Carl Gutwin, Professor, Department of Computer Science (966-8646)

Md. Sami Uddin, Grad Student, Department of Computer Science

This consent form, a copy of which has been given to you, is only part of the process of informed consent. It should give you the basic idea of what the research is about and what your participation will involve. If you would like more detail about something mentioned here, or information not included here, please ask. Please take the time to read this form carefully and to understand any accompanying information.

This study is concerned with evaluating the performance and learnability **of new kind of landmark-based menu – Artificial Landmark - Small.**

The goal of the research is to **evaluate the performance and learnability of command selection – Artificial Landmark system and compare with three different versions.**

The session will require **60** minutes, during which you will be asked to **select some items (icons) using the Artificial Landmark** system **which will be shown in desktop monitors** in the Human-Computer Interaction Lab at the University of Saskatchewan.

At the end of the session, you will be given more information about the purpose and goals of the study, and there will be time for you to ask questions about the research. As a way of thanking you for your participation and to help compensate you for your time and any travel costs you may have incurred, you will receive a **$10** honorarium at the end of the session.

The data collected from this study will be used in articles for publication in journals and conference proceedings.

As one way of thanking you for your time, we will be pleased to make available to you a summary of the results of this study once they have been compiled (usually within two months). This summary will outline the research and discuss our findings and recommendations. This summary will be available on the HCI lab's website: **http://www.hci.usask.ca/**

All personal and identifying data will be kept confidential. Confidentiality will be preserved by using pseudonyms in any presentation of textual data in journals or at conferences. The informed consent form and all research data will be kept in a secure location under confidentiality in accordance with University policy for 5 years post publication. Do you have any questions about this aspect of the study?

You are free to withdraw from the study at any time without penalty and without losing any advertised benefits. Withdrawal from the study will not affect your academic status or your access to services at the university. If you withdraw, your data will be deleted from the study and destroyed. Your right to withdraw data from the study will apply until results have been disseminated, data has been pooled, etc. After this, it is possible that some form of research dissemination will have already occurred and it may not be possible to withdraw your data.

Your continued participation should be as informed as your initial consent, so you should feel free to ask for clarification or new information throughout your participation. If you have further questions concerning matters related to this research, please contact:

- Dr. Carl Gutwin, Professor, Dept. of Computer Science, (306) 966-8646, gutwin@cs.usask.ca

Your signature on this form indicates that you have understood to your satisfaction the information regarding participation in the research project and agree to participate as a participant. In no way does this waive your legal rights nor release the investigators, sponsors, or involved institutions from their legal and professional responsibilities. If you have further questions about this study or your rights as a participant, please contact:

- Dr. Carl Gutwin, Professor, Dept. of Computer Science, (306) 966-8646, gutwin@cs.usask.ca
- Research Ethics Office, University of Saskatchewan, (306) 966-2975 or toll free at 888-966-2975.

Participant's signature:____

Date:____________________

Investigator's signature:____

Date:____________________

A copy of this consent form has been given to you to keep for your records and reference. This research has the ethical approval of the Research Ethics Office at the University of Saskatchewan.

Consent form used for medium interfaces.

DEPARTMENT OF COMPUTER SCIENCE
UNIVERSITY OF SASKATCHEWAN
INFORMED CONSENT FORM

UNIVERSITY OF SASKATCHEWAN

Research Project: **Artificial Landmark Medium – Spring 2016**

Investigators: Dr. Carl Gutwin, Professor, Department of Computer Science (966-8646)
Md. Sami Uddin, Grad Student, Department of Computer Science

This consent form, a copy of which has been given to you, is only part of the process of informed consent. It should give you the basic idea of what the research is about and what your participation will involve. If you would like more detail about something mentioned here, or information not included here, please ask. Please take the time to read this form carefully and to understand any accompanying information.

This study is concerned with evaluating the performance and learnability **of new kind of landmark-based menu – Artificial Landmark - Medium.**

The goal of the research is to **evaluate the performance and learnability of command selection – Artificial Landmark Medium system and compare with three different versions.**

The session will require **60** minutes, during which you will be asked to **select some items (icons) using the Artificial Landmark Medium system which will be shown in desktop monitors** in the Human-Computer Interaction Lab at the University of Saskatchewan.

At the end of the session, you will be given more information about the purpose and goals of the study, and there will be time for you to ask questions about the research. As a way of thanking you for your participation and to help compensate you for your time and any travel costs you may have incurred, you will receive a **$10** honorarium at the end of the session.

The data collected from this study will be used in articles for publication in journals and conference proceedings.

As one way of thanking you for your time, we will be pleased to make available to you a summary of the results of this study once they have been compiled (usually within two months). This summary will outline the research and discuss our findings and recommendations. This summary will be available on the HCI lab's website: **http://www.hci.usask.ca/**

All personal and identifying data will be kept confidential. Confidentiality will be preserved by using pseudonyms in any presentation of textual data in journals or at conferences. The informed consent form and all research data will be kept in a secure location under confidentiality in accordance with University policy for 5 years post publication. Do you have any questions about this aspect of the study?

You are free to withdraw from the study at any time without penalty and without losing any advertised benefits. Withdrawal from the study will not affect your academic status or your access to services at the university. If you withdraw, your data will be deleted from the study and destroyed. Your right to withdraw data from the study will apply until results have been disseminated, data has been pooled, etc. After this, it is possible that some form of research dissemination will have already occurred and it may not be possible to withdraw your data.

Your continued participation should be as informed as your initial consent, so you should feel free to ask for clarification or new information throughout your participation. If you have further questions concerning matters related to this research, please contact:

- Dr. Carl Gutwin, Professor, Dept. of Computer Science, (306) 966-8646, gutwin@cs.usask.ca

Your signature on this form indicates that you have understood to your satisfaction the information regarding participation in the research project and agree to participate as a participant. In no way does this waive your legal rights nor release the investigators, sponsors, or involved institutions from their legal and professional responsibilities. If you have further questions about this study or your rights as a participant, please contact:

- Dr. Carl Gutwin, Professor, Dept. of Computer Science, (306) 966-8646, gutwin@cs.usask.ca
- Research Ethics Office, University of Saskatchewan, (306) 966-2975 or toll free at 888-966-2975.

Participant's signature:____

Date:____________________

Investigator's signature:____

Date:____________________

A copy of this consent form has been given to you to keep for your records and reference. This research has the ethical approval of the Research Ethics Office at the University of Saskatchewan.

Consent form used for large interfaces.

DEPARTMENT OF COMPUTER SCIENCE
UNIVERSITY OF SASKATCHEWAN
INFORMED CONSENT FORM

UNIVERSITY OF SASKATCHEWAN

Research Project:	**Artificial Landmark Large – Summer 2016**
Investigators:	Dr. Carl Gutwin, Professor, Department of Computer Science (966-8646)
	Md. Sami Uddin, Grad Student, Department of Computer Science

This consent form, a copy of which has been given to you, is only part of the process of informed consent. It should give you the basic idea of what the research is about and what your participation will involve. If you would like more detail about something mentioned here, or information not included here, please ask. Please take the time to read this form carefully and to understand any accompanying information.

This study is concerned with evaluating the performance and learnability **of new kind of landmark-based menu – Artificial Landmark - Large.**

The goal of the research is to **evaluate the performance and learnability of command selection – Artificial Landmark Large system and compare with three different versions.**

The session will require **60** minutes, during which you will be asked to **select some items (icons) using the Artificial Landmark Large system which will be shown in desktop monitors** in the Human-Computer Interaction Lab at the University of Saskatchewan.

At the end of the session, you will be given more information about the purpose and goals of the study, and there will be time for you to ask questions about the research. As a way of thanking you for your participation and to help compensate you for your time and any travel costs you may have incurred, you will receive a **$10** honorarium at the end of the session.

The data collected from this study will be used in articles for publication in journals and conference proceedings.

As one way of thanking you for your time, we will be pleased to make available to you a summary of the results of this study once they have been compiled (usually within two months). This summary will outline the research and discuss our findings and recommendations. This summary will be available on the HCI lab's website: **http://www.hci.usask.ca/**

All personal and identifying data will be kept confidential. Confidentiality will be preserved by using pseudonyms in any presentation of textual data in journals or at conferences. The informed consent form and all research data will be kept in a secure location under confidentiality in accordance with University policy for 5 years post publication. Do you have any questions about this aspect of the study?

You are free to withdraw from the study at any time without penalty and without losing any advertised benefits. Withdrawal from the study will not affect your academic status or your access to services at the university. If you withdraw, your data will be deleted from the study and destroyed. Your right to withdraw data from the study will apply until results have been disseminated, data has been pooled, etc. After this, it is possible that some form of research dissemination will have already occurred and it may not be possible to withdraw your data.

Your continued participation should be as informed as your initial consent, so you should feel free to ask for clarification or new information throughout your participation. If you have further questions concerning matters related to this research, please contact:

- Dr. Carl Gutwin, Professor, Dept. of Computer Science, (306) 966-8646, gutwin@cs.usask.ca

Your signature on this form indicates that you have understood to your satisfaction the information regarding participation in the research project and agree to participate as a participant. In no way does this waive your legal rights nor release the investigators, sponsors, or involved institutions from their legal and professional responsibilities. If you have further questions about this study or your rights as a participant, please contact:

- Dr. Carl Gutwin, Professor, Dept. of Computer Science, (306) 966-8646, gutwin@cs.usask.ca
- Research Ethics Office, University of Saskatchewan, (306) 966-2975 or toll free at 888-966-2975.

Participant's signature:____

Date:____________________

Investigator's signature:____

Date:____________________

A copy of this consent form has been given to you to keep for your records and reference. This research has the ethical approval of the Research Ethics Office at the University of Saskatchewan.

B.2 QUESTIONNAIRES

Demographics and preference questionnaires.

Questionnaire - Overall

Sign in to Google to save your progress. Learn more

*Required

Participant ID *

Ask the experimenter if not provided

Your answer

Age *

In years

Your answer

Gender *

Select one.

- ○ Male
- ○ Female

Preference

Based on you experience from using three techniques, please answer the following questions.

Speed *

Which tech technique helped you to perform faster?

- ◯ Grid only
- ◯ Grid with Black-Blocks
- ◯ Grid with Image
- ◯ No Preference

Speed - Comment *

How did your preferred technique help you to perform faster?

Your answer

Accuracy *

Which technique helped you to perform your task accurately?

- ◯ Grid only
- ◯ Grid with Black-Blocks
- ◯ Grid with Image
- ◯ No Preference

Accuracy - Comment *

How did your preferred technique help you to perform accurately?

Your answer

Memorization *

Which technique helped you most to learn the position of icons easily?

- Grid only
- Grid with Black-Blocks
- Grid with Image
- No Preference

Memorization - Comment *

How did your preferred technique help you to memorize target locations easily?

Your answer

Expert Mode *

Which technique allowed you to use Expert Mode easily?

- Grid only
- Grid with Black-Blocks
- Grid with Image
- No Preference

Comfort *

Which technique did you find more comfortable?

- ○ Grid only
- ○ Grid with Black-Blocks
- ○ Grid with Image
- ○ No Preference

Overall *

Which technique do you prefer among the following three?

- ○ Grid only
- ○ Grid with Black-Blocks
- ○ Grid with Image
- ○ No Preference

Overall - Comment *

How did your preferred technique help you to memorize target locations easily?

Your answer

Submit   Page 1 of 1   Clear form

Never submit passwords through Google Forms.

NASA Task Load Index questionnaire.

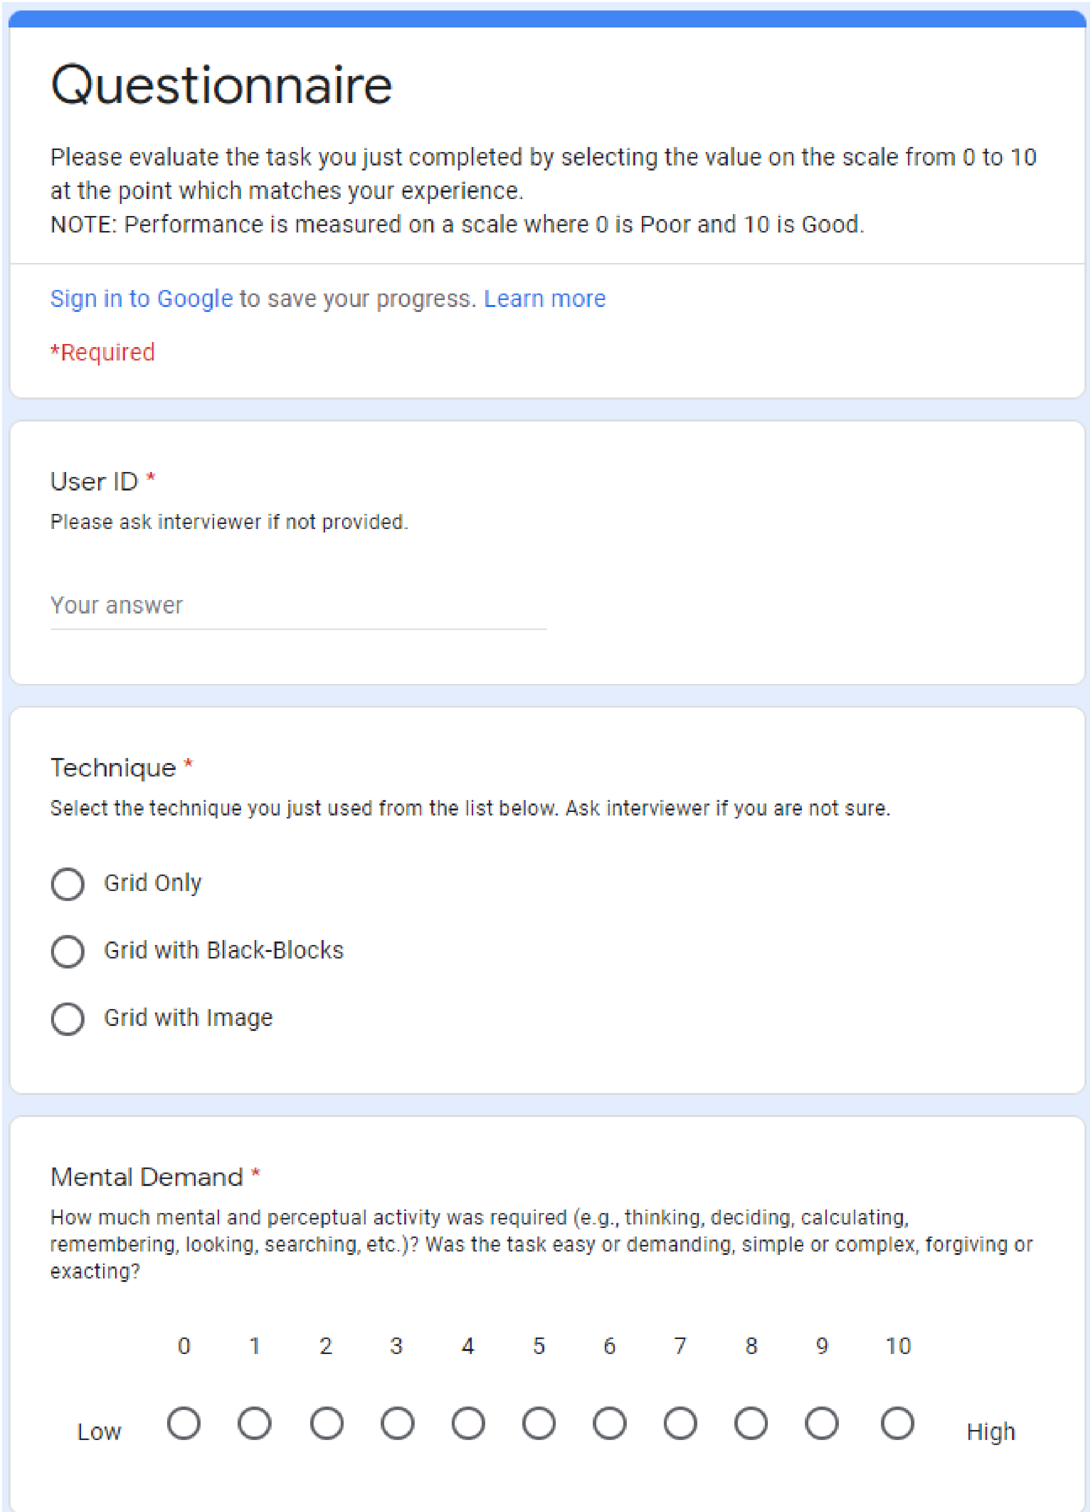

Questionnaire

Please evaluate the task you just completed by selecting the value on the scale from 0 to 10 at the point which matches your experience.
NOTE: Performance is measured on a scale where 0 is Poor and 10 is Good.

Sign in to Google to save your progress. Learn more

*Required

User ID *

Please ask interviewer if not provided.

Your answer

Technique *

Select the technique you just used from the list below. Ask interviewer if you are not sure.

- Grid Only
- Grid with Black-Blocks
- Grid with Image

Mental Demand *

How much mental and perceptual activity was required (e.g., thinking, deciding, calculating, remembering, looking, searching, etc.)? Was the task easy or demanding, simple or complex, forgiving or exacting?

	0	1	2	3	4	5	6	7	8	9	10	
Low	○	○	○	○	○	○	○	○	○	○	○	High

Physical Demand *

How much physical activity was required (e.g., pushing, pulling, turning, controlling, activating, etc.)? Was the task easy or demanding, slow or brisk, slack or strenuous, restful or laborious?

	0	1	2	3	4	5	6	7	8	9	10	
Low	○	○	○	○	○	○	○	○	○	○	○	High

Temporal Demand *

How much time pressure did you feel due to the rate at which the task elements occurred? Was the pace slow and leisurely or rapid and frantic?

	0	1	2	3	4	5	6	7	8	9	10	
Low	○	○	○	○	○	○	○	○	○	○	○	High

Performance *

How successful do you think you were in accomplishing the goals of the task set by the experimenter? How satisfied were you with your performance in accomplishing these goals?

	0	1	2	3	4	5	6	7	8	9	10	
Poor	○	○	○	○	○	○	○	○	○	○	○	Good

Effort *

How hard did you have to work (mentally and physically) to accomplish your level of performance?

	0	1	2	3	4	5	6	7	8	9	10	
Low	○	○	○	○	○	○	○	○	○	○	○	High

Frustration *

How insecure, discouraged, irritated, stressed and annoyed versus secure, gratified, content, relaxed and complacent did you feel during the task?

	0	1	2	3	4	5	6	7	8	9	10	
Low	○	○	○	○	○	○	○	○	○	○	○	High

Submit Page 1 of 1 Clear form

Never submit passwords through Google Forms.

APPENDIX C FOR MANUSCRIPT C

C.1 STUDY CONSENT FORMS

DEPARTMENT OF COMPUTER SCIENCE
UNIVERSITY OF SASKATCHEWAN
INFORMED CONSENT FORM

UNIVERSITY OF SASKATCHEWAN

Research Project:	**Artificial Landmarks – Winter 2017**
Investigators:	Dr. Carl Gutwin, Professor, Department of Computer Science (966-8646)
	Md. Sami Uddin, Grad Student, Department of Computer Science

This consent form, a copy of which has been given to you, is only part of the process of informed consent. It should give you the basic idea of what the research is about and what your participation will involve. If you would like more detail about something mentioned here, or information not included here, please ask. Please take the time to read this form carefully and to understand any accompanying information.

This study is concerned with evaluating the performance and learnability **of new kind of landmark-based media player and document viewer.**

The goal of the research is to **evaluate the performance and learnability of revisitation in media player and document viewer.**

The session will require about **30** minutes, during which you will be asked to **revisit specific location of the video and document using the specified interfaces shown in a desktop system** in the Human-Computer Interaction Lab at the University of Saskatchewan.

At the end of the session, you will be given more information about the purpose and goals of the study, and there will be time for you to ask questions about the research. As a way of thanking you for your participation and to help compensate you for your time and any travel costs you may have incurred, you will receive a **$5** honorarium at the end of the session.

The data collected from this study will be used in articles for publication in journals and conference proceedings.

As one way of thanking you for your time, we will be pleased to make available to you a summary of the results of this study once they have been compiled (usually within two months). This summary will outline the research and discuss our findings and recommendations. This summary will be available on the HCI lab's website: **http://www.hci.usask.ca/**

All personal and identifying data will be kept confidential. Confidentiality will be preserved by using pseudonyms in any presentation of textual data in journals or at conferences. The informed consent form and all research data will be kept in a secure location under confidentiality in accordance with University policy for 5 years post publication. Do you have any questions about this aspect of the study?

You are free to withdraw from the study at any time without penalty and without losing any advertised benefits. Withdrawal from the study will not affect your academic status or your access to services at the university. If you withdraw, your data will be deleted from the study and destroyed. Your right to withdraw data from the study will apply until results have been disseminated, data has been pooled, etc. After this, it is possible that some form of research dissemination will have already occurred and it may not be possible to withdraw your data.

Your continued participation should be as informed as your initial consent, so you should feel free to ask for clarification or new information throughout your participation. If you have further questions concerning matters related to this research, please contact:

- Dr. Carl Gutwin, Professor, Dept. of Computer Science, (306) 966-8646, gutwin@cs.usask.ca

Your signature on this form indicates that you have understood to your satisfaction the information regarding participation in the research project and agree to participate as a participant. In no way does this waive your legal rights nor release the investigators, sponsors, or involved institutions from their legal and professional responsibilities. If you have further questions about this study or your rights as a participant, please contact:

- Dr. Carl Gutwin, Professor, Dept. of Computer Science, (306) 966-8646, gutwin@cs.usask.ca
- Research Ethics Office, University of Saskatchewan, (306) 966-2975 or toll free at 888-966-2975.

Participant's signature:____

Date:__________________

Investigator's signature:____

Date:__________________

A copy of this consent form has been given to you to keep for your records and reference. This research has the ethical approval of the Research Ethics Office at the University of Saskatchewan.

C.2 QUESTIONNAIRES

Demographics and preference questionnaires.

Landmark Study - Overall

Please fill out the following questions carefully.

Sign in to Google to save your progress. Learn more

*Required

Participant ID *

Ask the experimenter if not provided

Your answer

Age *

In years

Your answer

Gender *

Select one.

- ○ Male
- ○ Female

Media Players *

What media player(s)/application(s) do you use regularly? E.g., vlc, windows media player, YouTube, Netflix, etc.

Your answer

MP Usage *

How much time do you spend on using media player in a week? Include web-based player (e.g., YouTube, Netflix, etc.) also.

- ○ 0 Hours
- ○ 1 - 10 Hours
- ○ 11 - 20 Hours
- ○ 20+ Hours

Document Viewers *

What document viewer(s)/application(s) do you use regularly to view (pdf) documents ? E.g., Adobe Acrobat reader, Foxit reader, etc.

Your answer

DV Usage *

How much time do you spend on using Document Viewer in a week (e.g., Adobe Acrobat reader, Foxit reader, etc.)?

- ○ 0 Hours
- ○ 1 - 10 Hours
- ○ 11 - 20 Hours
- ○ 20+ Hours

Next   Page 1 of 2   Clear form

Never submit passwords through Google Forms.

Preference

Based on you experience of using two systems, please answer the following questions carefully.

Landmarks *

What type of landmark was available in each of the two systems you just used? [Ask experimenter if you are not sure.]

	No Landmark	Random Icons	Thumbnails
Media Player	○	○	○
Pdf Viewer	○	○	○

Easiness *

How easy were the contents to remember?

	Very Easy	Easy	Normal	Hard	Very Hard
Video	○	○	○	○	○
Pdf	○	○	○	○	○

Video - Reason *

How did you remember the locations of the video? Did any feature of the player help you to revisit the video locations? If yes, please mention how?

Your answer

Pdf - Reason *

How did you remember the locations of the Pdf document? Did any feature of the Pdf Viewer help you to revisit the document locations? If yes, please mention how?

Your answer

Preferences *

Between the two systems which one would you prefer based on your performance for the following categories?

	Media Player	Pdf Viewer	None
Speed	○	○	○
Accuracy	○	○	○
Memorization	○	○	○
Comfort	○	○	○
Overall	○	○	○

Back Submit Page 2 of 2 Clear form

Never submit passwords through Google Forms.

NASA Task Load Index questionnaire used for media player.

Landmark Media Player Study

Please evaluate the task you just completed by carefully selecting the value on the scale from 0 to 10 at the point which matches your experience.
NOTE: Performance is measured on a scale where 0 is Poor and 10 is Good.

Sign in to Google to save your progress. Learn more

*Required

Participant ID *

Please ask the experimenter if not provided.

Your answer

Landmark Type *

What type of landmark was available in the slider of the player? [Ask the experimenter if you are not sure.]

Choose

Mental Demand *

How much mental and perceptual activity was required (e.g., thinking, deciding, calculating, remembering, looking, searching, etc.) to perform the task? Was the task easy or demanding, simple or complex, forgiving or exacting?

	0	1	2	3	4	5	6	7	8	9	10	
Low	○	○	○	○	○	○	○	○	○	○	○	High

Physical Demand *

How much physical activity was required (e.g., pressing, finger movement, controlling, activating, etc.) to perform the task? Was the task easy or demanding, slow or brisk, slack or strenuous, restful or laborious?

	0	1	2	3	4	5	6	7	8	9	10	
Low	○	○	○	○	○	○	○	○	○	○	○	High

Temporal Demand *

How much time pressure did you feel due to the rate at which the task elements occurred? Was the pace slow and leisurely or rapid and frantic?

	0	1	2	3	4	5	6	7	8	9	10	
Low	○	○	○	○	○	○	○	○	○	○	○	High

Performance *

How successful do you think you were in accomplishing the goals of the task set by the experimenter? How satisfied were you with your performance in accomplishing these goals?

	0	1	2	3	4	5	6	7	8	9	10	
Poor	○	○	○	○	○	○	○	○	○	○	○	Good

Effort *

How hard did you have to work (mentally and physically) to accomplish your level of performance?

	0	1	2	3	4	5	6	7	8	9	10	
Low	○	○	○	○	○	○	○	○	○	○	○	High

Frustration *

How insecure, discouraged, irritated, stressed and annoyed versus secure, gratified, content, relaxed and complacent did you feel during the task?

	0	1	2	3	4	5	6	7	8	9	10	
Low	○	○	○	○	○	○	○	○	○	○	○	High

Submit   Page 1 of 1   Clear form

Never submit passwords through Google Forms.

NASA Task Load Index questionnaire used for media player.

Landmark Scrollbar Study

Please evaluate the task you just completed by carefully selecting the value on the scale from 0 to 10 at the point which matches your experience.
NOTE: Performance is measured on a scale where 0 is Poor and 10 is Good.

Sign in to Google to save your progress. Learn more

*Required

Participant ID *

Please ask the experimenter if not provided.

Your answer

Landmark Type *

What type of landmark was available in the scrollbar of the Pdf viewer? [Ask the experimenter if you are not sure.]

Choose

Mental Demand *

How much mental and perceptual activity was required (e.g., thinking, deciding, calculating, remembering, looking, searching, etc.) to perform the task? Was the task easy or demanding, simple or complex, forgiving or exacting?

	0	1	2	3	4	5	6	7	8	9	10	
Low	○	○	○	○	○	○	○	○	○	○	○	High

Physical Demand *

How much physical activity was required (e.g., pressing, finger movement, controlling, activating, etc.) to perform the task? Was the task easy or demanding, slow or brisk, slack or strenuous, restful or laborious?

	0	1	2	3	4	5	6	7	8	9	10	
Low	○	○	○	○	○	○	○	○	○	○	○	High

Temporal Demand *

How much time pressure did you feel due to the rate at which the task elements occurred? Was the pace slow and leisurely or rapid and frantic?

	0	1	2	3	4	5	6	7	8	9	10	
Low	○	○	○	○	○	○	○	○	○	○	○	High

Performance *

How successful do you think you were in accomplishing the goals of the task set by the experimenter? How satisfied were you with your performance in accomplishing these goals?

	0	1	2	3	4	5	6	7	8	9	10	
Poor	○	○	○	○	○	○	○	○	○	○	○	Good

Effort *

How hard did you have to work (mentally and physically) to accomplish your level of performance?

	0	1	2	3	4	5	6	7	8	9	10	
Low	○	○	○	○	○	○	○	○	○	○	○	High

Frustration *

How insecure, discouraged, irritated, stressed and annoyed versus secure, gratified, content, relaxed and complacent did you feel during the task?

	0	1	2	3	4	5	6	7	8	9	10	
Low	○	○	○	○	○	○	○	○	○	○	○	High

Submit

Page 1 of 1

Clear form

Never submit passwords through Google Forms.