Space to Think: Sensemaking and Large, High-Resolution Displays

Christopher P. Andrews

Dissertation submitted to the Faculty of the
Virginia Polytechnic Institute and State University
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy
in
Computer Science and Applications

Christopher L. North, Chair
Douglas Andrew Bowman
Francis Quek
Narendran Ramakrishnan
John T. Stasko

August 9, 2011
Blacksburg, Virginia

Keywords: large displays, space, sensemaking, physical navigation, embodiment

Space to Think: Sensemaking and Large, High-Resolution Displays

Christopher P. Andrews

ABSTRACT

Display technology has developed significantly over the last decade, and it is becoming increasingly feasible to construct large, high-resolution displays. Prior work has shown a number of key performance advantages for these displays that can largely be attributed to the replacement of virtual navigation (e.g., panning and zooming) with physical navigation (e.g., moving, turning, glancing). This research moves beyond the question of performance or efficiency and examines ways in which the large, high-resolution display can support the cognitive demanding task of sensemaking.

The core contribution of this work is to show that the physical properties of large, high-resolution displays create a fundamentally different environment from conventional displays, one that is inherently spatial, resulting in support for a greater range of embodied resources. To support this, we describe a series of studies that examined the process of sensemaking on one of these displays. These studies illustrate how the display becomes a cognitive partner of the the analyst, encouraging the use of the space for the externalization of the analyst's thought process or findings. We particularly highlight how the flexibility of the space supports the use of incremental formalism, a process of gradually structuring information as understanding grows.

Building on these observations, we have developed a new sensemaking environment called Analyst's Workspace (AW), which makes use of a large, high-resolution display as a core component of its design. The primary goal of AW is to provide an environment that unifies the activities of foraging and synthesis into a single investigative thread. AW addresses this goal through the use of an integrated spatial environment in which full text documents serve as primary sources of information, investigative tools for pursuing leads, and sensemaking artifacts that can be arranged in the space to encode information about relationships between events and entities. This work also provides a collection of design principles that fell out of the development of AW, and that we hope can guide future development of analytic tools on large, high-resolution displays.

This research is partially supported by the National Science Foundation grants CCF-0937133, CNS-04-23611 and IIS-0736055, by the National Geospatial-Intelligence Agency contract HM1582- 05-1-2001, the National Visual Analytics Center (NVAC) and Information Infrastructure Integrity Initiative (I4) at Pacific Northwest National Laboratories, and the Institute for Critical Technology and Applied Science (ICTAS), Virginia Tech.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

In the last decade, there have been large changes in display technology. Display size and resolution have been growing, and multi-monitor systems are becoming more popular. These trends make the development of large, high-resolution displays increasingly more feasible. The term 'large, high-resolution display' can be defined in a number of ways. First, it can be defined in terms of the technology being used (i.e., 'larger than a traditional display'). This definition is subjective, as it implies that one's perception of the display itself is what defines it. As technology advances over time, and a 'traditional display' changes, displays that were once considered large (and high resolution) by this definition may no longer be. A second way to define the term is in terms of quantity of data that it can visually represent, or, perhaps more importantly, the ability to represent multiple views, scales, and 'units' of data (e.g., documents, web pages, etc.). However, this definition requires that we define what a unit of data is, which changes for each application or dataset.

Although both of these definitions provide useful points of reference, we prefer to define the term 'large, high-resolution display' as being a display that is *human scale*. By human scale,

we mean that the display's size and resolution are closely matched to the sphere of perception and influence of the human body. In a practical sense, the label describes displays that have a combined size and resolution that approach or exceed the visual acuity of the user. Displays at this scale afford the user the opportunity to trade *virtual navigation* (e.g., zooming, panning) for *physical navigation* (turning, leaning, moving around) [Ball et al. 2007], and thus allow the user to exploit embodied human abilities such as visual perception, spatial awareness, proprioception, and spatial memory. This tipping point is important because it heralds a change in user behavior, requiring new design considerations that are based on the extents of human abilities rather than the technological limitations of the display medium.

A number of studies have explored various benefits of using large, high-resolution displays (e.g., [Czerwinski et al. 2002; 2003, Ball et al. 2007, Yost et al. 2007, Kang and Stasko 2008, Shupp et al. 2009, Bi and Balakrishnan 2009]). However, there has been little consideration of the cognitive changes wrought by these displays. We believe that large, high-resolution displays are more than higher-bandwidth information conduits. Instead, the displays create an environment with the potential to transform tasks. In order to better explore the cognitive affordances of the display, our focus is on how these displays can be used for the cognitively demanding task of sensemaking.

1.0.1 Sensemaking

Sensemaking is the process of building understanding out of a collection of data. The process is often complex and ill-defined, involving data that is incomplete, dynamic and in some cases even wrong or deceptive [Heuer 1999, Thomas and Cook 2005]. It is involved in problems that include simple or regular activities like choosing a new phone or doing task management to more critical problems such as deconstructing what happened to the market or detecting

and stopping a terrorist plot before it happens. Sensemaking is a fundamentally human activity. Technology can provide support for searching, filtering, isolating, visualizing, and even identifying potential connections, but it cannot provide understanding. It is left to the human to conceptualize – to use judgment and intuition to identify the important, make logical connections, and draw conclusions. One of our great challenges as a research community is to develop tools and techniques that *join*, rather than replace, the human agent in the sensemaking process.

While sensemaking occurs in many domains, we focus primarily on intelligence analysis. Intelligence analysis has a couple of critical features as a domain of study. First, intelligence analysis is primarily concerned with human actions and motivations, which makes it less amenable to computationally driven analysis then some other areas, such as financial analysis or cyber security. Intelligence analysis is also a specialized form of knowledge work involving many of the same activities such as research, analysis, and knowledge management, making it easier to generalize our results to a wide base of potential beneficiaries. The difficulty of intelligence analysis has also made it the focus of study in an effort to understand the cognitive processes that underlie it (e.g., [Russell et al. 1993, Heuer 1999, Johnston 2003, Pirolli and Card 2005]), providing us with a foundation from which to form our theories about how large displays can become part of the process at a cognitive level.

1.1 Research Overview

The central thesis of this work is that *the physicality of large, high-resolution displays provides a fundamentally different environment than conventional displays, one that is inherently spatial, resulting in an alternate set of cognitive affordances that can be leveraged to aid in the sensemaking process.* This work is thus driven by three core research questions:

- What cognitive affordances are provided by large, high-resolution displays and how can they change the way users work and think?
- What advantages, if any, do large, high-resolution displays provide for the sensemaking process?
- How can we design analytic tools that leverage these advantages?

There are three major components to this work. The first is a pair of observational studies we performed to determine the potential role of large, high-resolution displays in the sensemaking process. These studies confirmed that the large display was viewed by subjects as a spatial environment, and demonstrated how the space was leveraged to aid memory and support organization-based synthesis. The second component was a study that explicitly compared the use of the physical space on a large, high-resolution display to a virtual workspace for a sensemaking task. This study showed that the greater support for embodiment provided by the large display led to more effective externalization of the sensemaking process into the available space. The third component is the development of a new analytic tool called Analyst's Workspace (AW), which was designed explicitly for use on the large, high-resolution display, and is based on the findings of the studies.

1.1.1 Theoretic Approach

The underlying perspective that informs this work is one of distributed cognition. The central tenet of distributed cognition is that people form tightly coupled cognitive systems in concert with their environments [Hollan et al. 2000, Hutchins 2001]. As such, humans recruit external resources as part of the problem-solving process, transforming the problem into one requiring a different set of cognitive abilities [Clark 2008, Hutchins 2001]. In other

words, humans are *embodied* beings, and external abilities developed for interacting with the outside environment such as perception and locomotion are as fundamentally part of the cognitive process as internal systems such as memory [Gray and Fu 2004].

This perspective is particularly useful for this work for two primary reasons. First, it provides initial insight into the role of the display. The process of sensemaking is inherently about the creation of understanding, which as stated above, is a fundamentally human process. As such, the role of the display is to support the human analyst, bridging the gap between human cognitive abilities and information and computational resources. The goal is for the analyst to be able to work through the display and interact directly with the data. Embodied interaction tells us that this should come through the creation of a wide range of ways to interact with the system, ways that make use of a broad spectrum of embodied resources [Dourish 2001]. In addition, if we consider the analyst and environment as a coupled system, there is an implication of bidirectional information flow [Clark 2008]. If the environment is truly part of the sensemaking process, it is not merely a producer of information that the analyst consumes, it needs to be an active part of not just foraging for information, but also the act of synthesis, where that information is shaped into a coherent narrative. In other words, it is important that the system supports the analyst in externalizing information back out into the environment where he or she has the opportunity to interact with it.

The second useful feature of distributed cognition is that it provides guidance in our observation studies and analysis as we attempt to determine the utility of the display for sensemaking. If the display and analyst are a coupled cognitive system, it can be studied by observing the interchange of external representations [Zhang 1997, Hutchins 2001]. This leads to questions such as:

- What external representations are created?

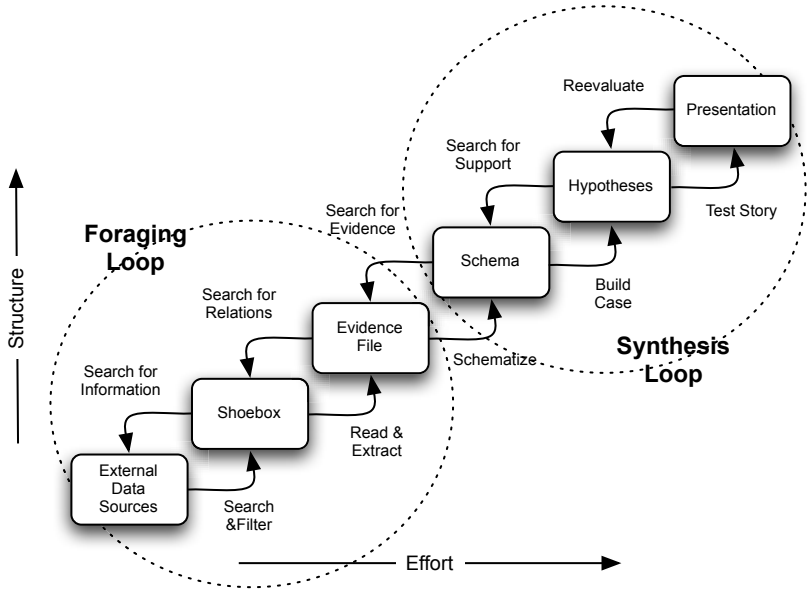

Figure 1.1: The sensemaking loop, adapted from [Pirolli and Card 2005]

- How are they created?
- How do they evolve?
- How does the evolution of the representations transform into understanding? [Liu et al. 2008]

For our representation-based analysis of the sensemaking process, we turn primarily to Pirolli and Card‘s cognitive model of how sensemaking works [Pirolli and Card 2005]. Their model of the “sensemaking loop“ breaks the sensemaking process down into a series of distinct phases based on a combination of the action performed by the analyst and the class of artifacts generated by the action (Figure 1.1). The sensemaking loop helps us to derive the meaning of a representation and identify where it fits within the whole of the sensemaking process. Pirolli and Card have also identified a number of “leverage points“ in this process Ð problematic points in the process that would benefit from improvements. These are primarily concerned with issues of data overload and attention management. Our explorations of the characteristics of large, high-resolution displays directly addresses some of these problem

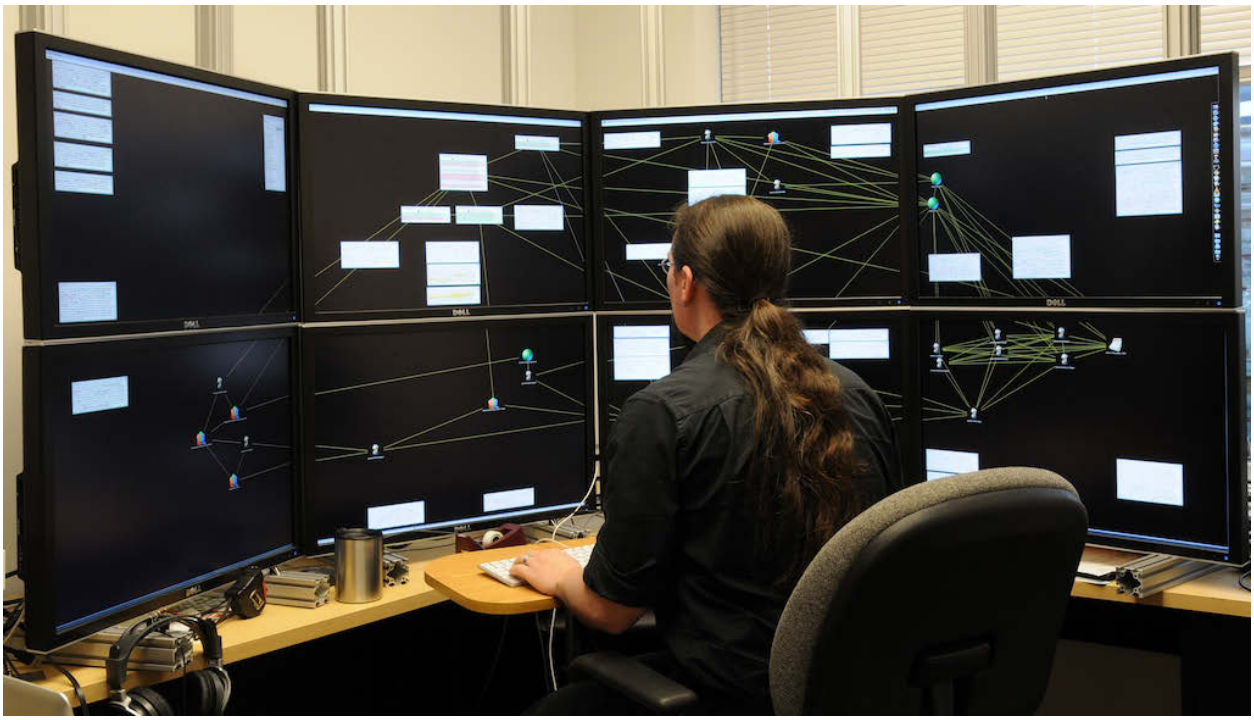

Figure 1.2: The 33 megapixel large, high-resolution display running Analyst's Workspace.

areas.

1.1.2 The Display

While our definition of large, high-resolution displays covers a considerable range of possible implementations and forms, all of the work described here was conducted using displays with identical form factors. The display is a 4x2 grid of 30" monitors, each with resolution of 2560x1600 pixels (Figure 1.2). This creates a display with an overall resolution of 10240x3200 or about 33 megapixels. Each column is supported by a freestanding support that allows the display to be reconfigured to curve around the user, which has been shown to lead to more efficient use [Shupp et al. 2009].

There are a couple of practical benefits for using this display. The pixel density of the 30" monitors means that we get a large number of pixels with a relatively small footprint, making

it practical to put the display in an office environment, where it could be used all day, rather than having to install it in a special lab setting. It can also be driven by a single machine, which means that conventional desktop applications can be run without modification and common development tools and libraries can also be used, making development of new tools easier. The use of a single machine also allows us to use the standard mouse and keyboard for interaction. As can be seen in the figure, we use a rolling podium for the input devices so that the user can move them to work on different locations on the display.

1.2 Organization

We begin our discussion in Chapter 2 with our theory of the spatial nature of large, high-resolution displays and potential uses for sensemaking. This chapter surveys the state of the field and lays out the basic theoretic framework that was developed in the course of the studies. Chapter 3 describes the pair of studies we conducted to explore the potential of the display itself as a sensemaking tool. This work informed the spatial approach that we adopted and has been published previously Andrews et al. [2010]. In Chapter 4 we discuss the third study that explicitly compares the large, high resolution display to a virtual workspace for sensemaking based on spatial organization of information. Chapter 5 describes the development of Analyst‘s Workspace, and finally in Chapter 6 we summarize the work and contributions.

Chapter 2

Foundations and Related work

2.1 Large, High-Resolution Displays

At their simplest, large, high-resolution displays allow users increased simultaneous access to information. A typical 17″ monitor covers only about 10% of the visual field and only about 1% of what can be seen allowing for head movement [Grudin 2001]. The impoverished environment of the single monitor forces users to make explicit context switches on the introduction of new information, frequently in the form of a new window overlaying the previous one. This in turn severely affects the user's ability to make comparisons and requires the user to expend valuable mental resources on the minutiae of managing views rather than on the problem at hand. Increasing the display size changes that dynamic, allowing the user to access more information at once. Comparisons can be done visually, rather than relying on memory and imperfect internal models. A flick of the eye or turn of the head is all that is required to consult a different data source. These advantages have been observed for even relatively small large, high resolution displays [Czerwinski et al. 2003, Consulting 2005, Kang

and Stasko 2008]. Importantly, Yost has shown that these benefits extend well beyond the limits of visual acuity, demonstrating that there is more to large, high-resolution displays than just covering more of the visual field [Yost et al. 2007].

2.1.1 Spatial environments

Increasing the display size and resolution changes more than merely the quantity of visually available information. We agree with Swaminathan and Sato when they state "We think that when a display exceeds a certain size, it becomes qualitatively different: different design issues come into play and interaction design becomes full-blown environment design" [Swaminathan and Sato 1997]. We would go a step further and say that the differences go beyond issues of design. We feel that there is a qualitative difference from a cognitive perspective as well – large, high-resolution displays are *perceptually* different. Based on the following studies, it is our contention that the quantity of information that can be displayed and the physical size and uniformity of the display work together to form a *spatial environment*, creating spatial relationships not just between the displayed objects, but also between the user and locations on the display.

Tan provides interesting evidence for this shift in a collection of studies he conducted that completely isolated display size as a factor [Tan 2004]. In his studies, the large display was a projected image at the same resolution as the conventional display, and the user was moved in relation to the display so that the resulting retinal image would be the same size. Interestingly, despite the fact that both displays contained the same information and covered the same viewing angle for the user, Tan still demonstrated performance advantages for the large display for 3D wayfinding and spatial understanding tasks. In further studies, Tan discovered that the large display was biasing users to take an egocentric view of the

displayed content, which was in turn aiding them in performing spatial tasks. This egocentric bias is important because the egocentric frame of reference is the primary axes around which location and layout is perceived and remembered [Biocca et al. 2001, Avraamides and Sofroniou 2006]. In addition, Patten and Ishii suggest that the ability to use an egocentric frame of reference creates opportunities to assign semantics to the space based on associations with the body [Patten and Ishii 2000].

Biasing users to adopt an egocentric frame of reference is not the only way in which large, high-resolution displays support spatial understanding. Czerwinski et al. performed a series of experiments which used a large $36'' \times 14''$ (2048×768) display to examine the effect of field of view (FOV) on basic 3D navigation tasks. The main focus of these experiments was to examine the how the change in the quantity of information would affect the disparity in spatial abilities between the genders. They found that the large display helped both genders to navigate the virtual worlds faster, and that women benefited more from the increased FOV than men. The underlying hypothesis is that since women depend more on landmarks for spatial navigation, they benefit more from the offloading of cognitive map-building to the perceptual system that is afforded by the increased available information [Czerwinski et al. 2002].

In a similar study, Ni et al. examined the roles of both resolution and size in performing simple spatial tasks such as navigation, search and comparison within an information-rich virtual environment (a VE that has been enriched with related abstract data). Both display size and resolution were varied, with the FOV determined by the display size (i.e., the larger display showed more spatial information, and the higher resolution showed better quality textual information). Both increased size and increased resolution had a positive effect on performance, but the effect of size was far greater than the effect of resolution [Ni et al. 2006].

In both of these studies, the subjects benefited from the increased field of view afforded by the larger display.

Another change made by large, high-resolution displays is an increase in physicality. Ball has performed several important studies that examined how the size of the display affects user performance and behavior. Subjects were presented with a geospatial visualization containing embedded information about houses for sale, and asked to perform a number of tasks that can be broadly categorized as search, navigation, pattern finding and insight generation. As the size of the display increased, performance on the search and navigation tasks improved. More critically, this experiment explicitly examined the tradeoff between virtual navigation and physical navigation. As the size of the display increased, virtual navigation was increasingly replaced by physical navigation and Ball showed a correlation between the amount of virtual navigation and performance (as virtual navigation decreased, performance increased) [Ball et al. 2007]. In a followup study, Ball used a similar experimental setup but isolated the effects of peripheral vision and physical navigation. He found that conditions allowing physical navigation positively affected performance, while peripheral vision had no significant effect on performance [Ball and North 2008]. Shupp built on these results and demonstrated that curving the display around the user increased performance even further, due to the increased efficiency of the physical navigation (i.e., it is quicker to turn than to move laterally) [Shupp et al. 2009]. Thus we find that as the size of the display increases, behavior changes and users begin to treat the displayed information as physically part of their environment, trading abstract system interaction with physical movement.

While all of these studies demonstrate positive benefits of large displays, Jakobsen and Hornbæk provide a note of caution. They performed a series of comparative studies very similar to those performed by Ball and Shupp that examined how display size affected the performance of geospatial tasks [Rønne Jakobsen and Hornbæk 2011]. However, they added

interaction and information spaces that were larger than the largest display, requiring the use of virtual navigation techniques in all conditions. They reported that while physical navigation did indeed increase, performance started to drop on the largest display due to the increased area that needed to be scanned for targets and interacted with. So, while this continues to support the idea that the physicality of the display resulted in an increase in the physicality of the interaction, it also suggests (as Swaminathan and Sato did) that conventional interaction design does not simply scale up [Swaminathan and Sato 1997].

Working with non-spatial information

The examples we have examined have been explicitly spatial - either geospatial visualizations or virtual environments in which space was a primary element. When Tan looked for benefits of the large display for tasks with no spatial component such as text comparison, he found no benefits [Tan 2004]. Similarly, Yost's work on the scalability of visualizations showed that spatially based visualizations (e.g., individual visualizations embedded in a map) scaled better than an attribute based visualization (e.g., small multiples), implying that developers should attempt to find ways to transform non-spatial tasks into spatial ones [Yost et al. 2007].

To address this, it is useful to consider the potential of the resolution to creating space in subtle ways based on the interaction between the size of a representational object and the size of the containing medium. The meaning that can be attached to an object's position is directly related to the amount of space consumed by its representation with relation to the total available space. On conventional displays, creating a spatial environment leads to the use of very small representations ranging from simple dots, to icons, labels and thumbnails. Each of these provides varying amounts of visually available information with associated

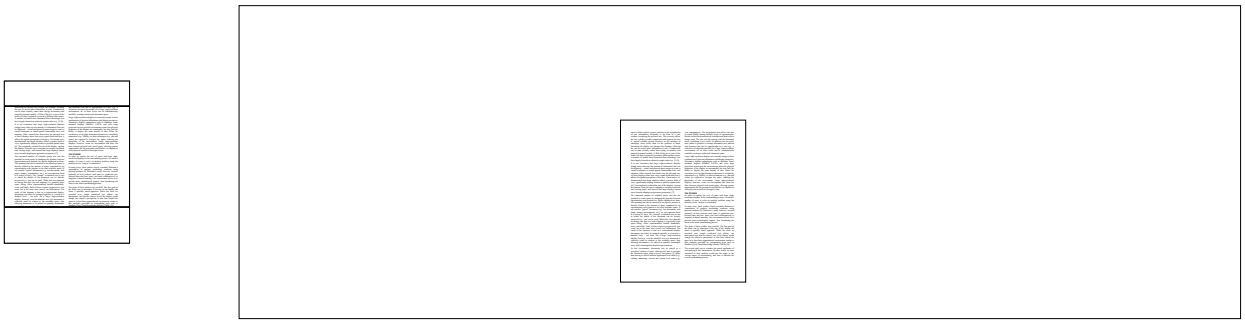

Figure 2.1: Creating space through the change in ratio between representation and display size - (left) a document maximized on a 1280×1024 display (right) the same document displayed at the same size on a 10240×3200 display

tradeoffs in the number of elements that can be displayed and meaningfully arranged. This tradeoff is a very important design consideration as a difference in milliseconds in the time required to access information may spell the difference in whether users consult the environment or use their own memory [Gray and Fu 2004]. Russell et al. provide a nice illustration of this in the study of their Grokker tool, a tool for visualizing and exploring a document collection. By redesigning the tool to represent documents using the first few words of the content, rather than an unlabeled dot with the start of the document available through a tooltip, they dramatically increased the amount of reading done by the users [Russell et al. 2006].

Large, high-resolution displays provide an opportunity to expand the range of document representations that are available to be used spatially to include what we might call the 'detailed view' or the 'working view'. This is the view in which the text of the document is fully available for reading or editing. Documents, of course, can vary widely in length, so we will consider the view to be about a 'page' in length, with other content available through scrolling or other navigation techniques. This is a compelling size because it is large enough for most news stories and well-designed web pages, is a comfortable length to read, matches the physical artifacts that we associate with most documents (i.e., a sheet of paper), and ignoring the minor differences in size between different formats (e.g., letter versus A4),

remains a standard unit length for even digital documents. On most standard monitors, even a single page cannot be displayed fully at a comfortable size for reading and editing. Large, high-resolution displays, however, can be sufficiently large such that the available display space is several times larger than the space required to comfortably display a single page, thus allowing documents to be placed in spatially meaningful ways while retaining their detailed representations (Figure 2.1). In other words, large, high-resolution displays have the potential to provide an experience akin to working on a real desktop, finally supplying the spatial dimension to the 'desktop' metaphor.[1]

Another way in which this ratio between representation and display size emulates space is through the support of multiscale representations. Daily interaction with the world makes use of the multiple levels of resolution and hierarchical nature of space [Kuhn 1996, Darken and Sibert 1996]. As a rough example, an observer standing in front of a building could be regarding the building, the door leading into the building, a sign posted on the door, a word on the sign, or a character contained within the word. All of these levels of resolution (and more) are simultaneously present in the world and available to the observer through a simple change in regard. Because of this ratio between representation size and display size, there are many potential opportunities for integrated or multiscale visualization. For example, one interpretation of placing a document in a "spatially meaningful" position could be to use it as an entity in a visualization such as a cluster or a timeline. By integrating full documents into a clustering scheme, the user can rapidly switch between regard for the entire space, individual clusters, or the contents of a document in the cluster, "zooming in" to look at details just by shifting regard. Considering Furnas and Bederson's Space-Scale diagram, the

[1]It is worth noting the potential effect of monitor bezels in tiled displays in supporting this dynamic. It is well established that users are typically uncomfortable with windows that straddle bezels when using multiple monitors [Grudin 2001, Hutchings et al. 2004]. Unsurprisingly, it has been our experience that this behavior also manifests itself on tiled large, high-resolution displays. With small tiles, users tend to fill each tile with a window, limiting the flexibility of the space. This makes it important that the individual tiles be large enough to display an entire document with some room to move.

physicality of the environment allows the user to traverse the pyramid of potential views of the available multiscale space through physical movement [Furnas and Bederson 1995]. Yost suggests that this form of integrated visualization is one of the best ways to take advantage of large, high-resolution displays [Yost et al. 2007].

Bi and Balakrishnan provide compelling evidence for the inherently spatial nature of large, high-resolution displays [Bi and Balakrishnan 2009]. They performed a longitudinal study of students working on a large, high-resolution display (a $16' \times 6'$ wall with a resolution of 6144×2034 made of tiled projections) for their normal work over the course of a week and compared it to how they worked in their normal work environment. One major distinction between the conditions was the way user's arranged their windows. For participants with dual monitors, they found little effort was made to arrange windows beyond the use of one display for primary work and the other for more peripheral information. On the large display, however, the participants spent far more effort trying to optimize the layout to improve workflow, with one participant reporting that he spent the first five minutes of every day laying out the space. The participants also tended to mentally partition the large display space into focal and peripheral regions, with most work being performed in the central focal region and secondary applications such as email and instant messaging clients being placed around the periphery of the display. Another important observation was that application windows rarely spanned more than a fraction of the large display. Only one participant ever maximized a window, and that was to show a map. Even when examining large, information rich documents such as spreadsheets, only 60-70% of the display was required. This observation is key, especially when taken in concert with the observation about the increased attention to layout. This demonstrates how the display has changed user perception and behavior, making space and spatial relationships an important part of their everyday working environment.

This section has demonstrated a number of ways in which the attributes of large, high-resolution displays contribute to a user perception of a spatial environment in which locations have meaning within the user's own egocentric frame of reference. In the next section we will examine how this environment can support the sensemaking process.

2.2 Space and Sensemaking

Space plays a key role in the human cognitive system due to our experience of living and acting in space [Kuhn 1996]. According to Lakoff and Johnson, "Most of our fundamental concepts are organized in terms of one or more spatialization metaphors" [Lakoff and Johnson 1980]. These metaphors can be used to create powerful affordances that leverage human perception and memory to reduce cognitive load, simplify choice, and transform tasks to more familiar forms [Kirsh 1995, Kuhn 1996]. In this work, we are primarily interested in how space can be used as part of the sensemaking process, in particular, the latter stages of evidence marshaling and synthesis which is particularly difficult and ill-supported. Despite this difficulty, it is still largely this part of the process that is done in the analyst's head [Wright et al. 2006]. By performing this difficult task internally, analysts are more prone to cognitive bias and reduce their ability to communication their hypotheses and the analytic process they followed to reach them [Heuer 1999]. In the course of the studies, we have distilled the use of the spatial environment afforded by large, high-resolution displays down to two broad, interconnected categories: memory and organization.

2.2.1 Memory

Many of the ways in which space can support the sensemaking process can be roughly classified as supporting memory. Space provides the potential for external representations

to be *persistent*. Given a spatial environment there is an opportunity to display information in a fixed location with a high level of detail, in other words, the information can have both spatial and representational persistence.

External representations are the primary unit of interaction and communication in distributed cognitive systems. They can act as mediators; their form structuring and changing the nature of tasks [Zhang and Norman 1994, Norman 1991, Hollan et al. 2000] . We focus here, however, on one of their most obvious and widely accepted roles, that of memory aid. One of the core concepts in embodiment is that rather than building rich, detailed inner models, we typically rely on repeated consultations with the outside world [Clark 1999]. O'Regan proposes that this is an explanation for the apparent discrepancy between how the world is visually perceived and what we know about the actual properties of the human eye. He posits that the world provides an "external memory" that is continually accessed to provide the seamless, high-fidelity, in focus visual perception that we know the eye is unable to provide as an isolated optical instrument [O'Regan 1992].

External representations are not limited to the range of immediately perceivable environment. They could include a file buried in a file system, a book in a library, or a piece of paper at the bottom of a stack on a desk. Their benefit for memory is immediately obvious in that they are a form of "perfect" memory (for themselves - consulting a string tied around a finger will reveal everything about the physical properties of the string and knot used to attach it, but not necessarily anything about any additional information intended to be associated with it). The use of external representations for memory does come with a cost over the use of internal memory - the cost of increased access time. Gray has demonstrated that humans employ "micro strategies" to optimize their actions at a very low level, and that a couple of milliseconds in extra access time can make the difference between whether an external representation is consulted or the less reliable internal memory is used [Gray

and Fu 2004]. One of the benefits of external representations that are spatially available and persistent is that the cost to access is relatively low. Physical navigation, be it turning, moving, or just glancing, is all that is required to consult the representation. Plumlee and Ware demonstrated this effect in their work on visual working memory. They found that for comparison tasks when interaction (zooming) was required to access the separate representations being compared, subjects made far fewer accesses than their model of visual working memory predicted, preferring, instead, to guess based on a limited number of shared features. When the two representations were both visually available, subjects made more accesses than the model predicted would be required due to the ease in which the accesses could be made [Plumlee and Ware 2006].

Spatial memory

There are two general ways to think about space and memory. Consulting external representations is a way to use space for memory. Spatial memory, on the other hand, is about actual location and addresses the questions "what is here?" and "where is?" [Kuhn 1996]. Richardson provides compelling evidence that the human brain treats location information specially and can even "tag" non-spatial information with location information [Richardson and Spivey 2000]. In his "Hollywood Squares" experiments, he presented information to subjects with some associated locational cue. He found through the use of eye trackers that when prompted for the information, subjects would consult the location associated with the information even when the screen was blank. Interestingly, this is not at a conscious level as the subjects were unaware of doing this and frequently could not actually say which location was associated with the information. This does provide a compelling argument for spatial persistence, however, since the perceptual system can rapidly take over after the initial eye movement.

Spatial memory has been put forth as a powerful cognitive tool for rapidly accessing information [Robertson et al. 1998]. It is difficult, however, to gauge the efficacy of conscious spatial memory in isolation. For example, Jones and Dumais showed that self-selected two letter mnemonics out-performed location encoding for article filing and recovery [Jones and Dumais 1986]. In considering this result, however, it is worth examining the nature of the task performed by the subjects. Subjects were given short news articles covering a broad spectrum of topics which once filed remained filed (i.e., unreadable). So the ability to relocate a file rested entirely on the initial choice of encoding at the time of filing, and there was little opportunity for location learning. In a followup study, they also found that subject performance with the two-letter mnemonics generated by other subjects was better than chance, suggesting that the performance of the mnemonics included some analytic component beyond pure recall.

Cockburn performed a number of studies looking at various attributes of spatial memory that illustrate some of the difficulty of separating spatial memory from other phenomena [Cockburn and McKenzie 2002, Cockburn 2004]. In one study, he asked subjects to remember the location of flags in a two dimensional space and found that subjects were generally very poor at this unless the flag they were looking for was one with which they were already familiar. This is unsurprising given the complexity of remembering the appearance of an arbitrary representation without being able to tie it to some higher level abstraction [Plumlee and Ware 2006]. In a second study, Cockburn found that subjects performed much better when the spatially encoded representations were letters, but only when they were arranged in rows and not columns. In questioning the subjects he found that in the horizontal case, they were chunking the letters and forming mnemonics to help them remember order [Cockburn 2004].

Ehert also provides an interesting study of this dynamic [Ehret 2002]. Ehert was primarily

concerned with the mechanics of location learning. In his study, he presented subjects with a color stimuli surrounded by buttons. Given a particular color, the subjects were required to click the corresponding button as fast as possible. There were several conditions ranging from colored buttons, textual buttons, to buttons with text or images that had no correspondence to the color. Unsurprisingly, the condition with the colored buttons performed the best. However, over a large number of trials, the results evened out with no advantage for any condition. In other words, the subjects learned the location of the buttons and no longer had to consult the encoding. Interestingly, in the second phase of the study, the screen was cleared and the subject was asked to recreate the configuration. For this task, the colored buttons performed the worst, suggesting that location was never learned, performance was driven purely by pre-attentive processing and recognition.

The method of loci is another example of the role of location in cognition that reverses the way we typically think about spatial memory, using location as mnemonic to recall sequential information rather than a way to locate objects physically in space [Hess et al. 1999]. This long used technique explicitly associates ordered information (e.g., lists, stories, songs, etc) with a well-learned space. The structure of the space provides the retrieval cues that help to trigger the next item in the sequence. Ragan et al. applied this concept in a study of the use of virtual environments for process memorization, and showed that richer spatial cues in the form of greater field of view and greater field of regard led to better recall [Ragan et al. 2010]. While the focus of this work was on virtual environments, the core finding that the greater spatial cues afforded by increased field of view and field of regard led to improved use of the space lends additional credence to our contention that the physical characteristics of large, high-resolution displays lead to an environment that provides greater opportunities to leverage human spatial abilities than an environment provided by a conventional display.

Across all of this research, we find evidence that location plays an important role in cognition.

However, it seems that it is most useful when the location is meaningful, either through prior connections or through learning and the creation of meaning. In addition, it seems that spatial memory is frequently coupled with perception and recognition, with users seamless transitioning between them without realizing that they have done so [Blanc-Brude and Scapin 2007]. This provides an argument for workspaces that maintain spatial and representational persistence or invariance.

External memory for sensemaking

External memory can play a number of roles in the sensemaking process. The most obvious form is as an extension of the analyst's working memory, allowing them to attend to more evidence and more complex structures [Pirolli and Card 2005]. The advantage of using the spatial environment is to speed up the access time and encourage reliance on the actual information rather than the imperfect internal memory of the analyst. While this takes advantage of the ability to rapidly access information in the space, a second advantage of space is the high availability of the information. In looking at how office workers arranged their desks, Malone found that desk organization was as much about reminding as it was about rapidly finding a particular document [Malone 1983], and in the spatial environment provided by a large, high-resolution display can provide similar affordances. A third use would be to contextualize information. Various invariant or well-learned attributes of a location could inform about potential implied relationships, or help an analyst regain a "state of work" awareness when returning to a task or document [Fass et al. 2002].

A less direct use of external memory is to use it to support other cognitive activities. We find a potential example of this when we reexamine Ball's studies of physical versus virtual navigation [Ball et al. 2007, Ball and North 2008]. Ball observed an increase in performance that

corresponded with the increase in physical navigation. This could be interpreted as a purely mechanical effect - the subjects found it faster to move than to virtually navigate and as the display size increased and with it the potential to physically navigate, they opportunistically shifted to physical navigation. While this may explain some of the effects observed in the study, it doesn't take into account the minimal effort required to navigate large distances virtually. A different interpretation of Ball's results is that the larger displays are also providing cognitive benefits that partially account for the performance improvements. Using a large display, the user has no need to maintain an internal model of the environment - the environment is its own model. Bederson found that adding animation to transitions through virtual space increased the user's spatial knowledge of the environment [Bederson and Boltman 1999]. This is immediately relatable to the results of Ball's study in which optical flow and proprioception help to maintain the coherency of the model and provide real physical cues as to the spatial nature of the data, and thus better supporting the subject's development of spatial understanding.

2.2.2 Organization

While spatial organization can be a form of external memory, it can also be part of more complex cognitive processes. In his study of how analysts synthesize the results of their analysis into meaningful conclusions, Robinson examined some of the ways that analysts can leverage this organizational property of space [Robinson 2008a]. Robinson created a synthetic dataset, which he provided to analysts as physical artifacts on notecards. The analysts worked with the information directly, organizing it on a large table and annotating the space. One of the key observations was that as the analysts worked with the information, they were arranging it spatially, using a variety of spatial metaphors, such as clustering and

timelines, to understand the information. In working with external representations in this fashion, the analysts were using space to create a *semantic layer*, adding meaning based on spatial relationships. This meaning is created not just through position, but also through the act of positioning. In other words, meaning can arise from the "discussion" implied by the arrangement and rearrangement of objects in space [Nakakoji et al. 2000].

Kirsh identifies three primary uses for space beyond its use as external memory: simplifying choice, simplifying perception, and saving internal computation [Kirsh 1995]. We can use space to simplify choice by creating cues and constraints in the environment. For example, when disassembling a complex piece of equipment, one might arrange the parts in space according to the order they were removed. When reassembling, the parts can be reattached in the correct order simply by traversing the space in the opposite direction. A simpler example might be the placement of implements on a stovetop. Placing each implement near the pot with which it is being used creates a visual and physical cue as to which tool to use and reduces the likelihood of cross-contamination between pots.

Simplifying perception by using space is one of the primary techniques of information visualization. For example, clustering makes categorizations more explicit, making it easier to identify not just what categories are present, but also information about the categories, such as the number of members or relative size. Adding some explicit semantics to the space and we get familiar visualizations, such as timelines and scatterplots, that can reveal properties of the data that were not previously obvious. However, as Nakakoji et al. point out, there are two ways to think about position: position as the state of an object, and the act of positioning [Nakakoji et al. 2000]. Positioning can take the form of "discussion" with the data, based on a tight feedback loop with the perceptual system. This kind of discussion is the root of data immersion, one of the fundamental strategies for performing analysis identified by Heuer. Data immersion is based on sorting and organizing information

until meaningful patterns emerge [Heuer 1999]. The use of this technique is essentially what Robinson observed in his studies [Robinson 2008a].

There are varying degrees to which space can used to save internal computation. For example, counting on one's fingers transforms a problem of mental arithmetic to one of physical representation and counting. Abacuses, slide rules, and calculators turn complex arithmetic into mechanical processes. A subtler example comes from Kirsh and Maglio's study of Tetris players [Kirsh and Maglio 1994]. They discovered that, contrary to intuition, expert Tetris players rotated pieces *more* than inexperienced players. Rather than seeing the next piece and computing the best fit, only rotating a piece enough to align it properly, expert players spun the pieces as they appeared, using perception to guide them to the best fit. Kirsh and Maglio refer to this behavior as epistemic action, action performed in aid of computation, as opposed to pragmatic action, action performed in the furtherance of a goal.

Incremental formalism

One advantage of using space to think is that it offers tremendous flexibility, allowing for experimental exploration. As Shipman points out, many tools require users to structure their data as they enter it, either to make it easier for the system itself to manage the information, or ostensibly to help users understand their information better through the nature of the structure. While there is merit in both of these aims, it is often the case that users are forced to perform this formalism too early. There is a high cognitive cost involved in trying to resolve the mismatch between the system language and a user's understanding of the information he or she is trying to impart to the system. Frequently, this results in the user being forced to prematurely commit to explicit structures, leading to incorrect or inconsistent formalization that may be of less value than completely unstructured information [Shipman

and Marshall 1999b]. The analytic process has been described as a dynamic, free-flowing cycle of hypothesis formation, testing and rejecting [Pike et al. 2007]. Early commitment to structure also inhibits this process, impeding the exploratory aspects of the process, not just forcing users to make uninformed choices, but also limiting change.

A particularly compelling example of the problems of early formalization is the experiences reported by the developers of Aquanet [Marshall and Rogers 1992]. Aquanet was designed to make it easier for users to express complex ideas by providing a robust collection of structures and structure building tools. In order to enter information into the system, users had to first specify a datatype to use to contain it. However, when given an exploratory knowledge building task, users often reported that it was difficult to know how they were going to use a piece of information when they were first entering it. As a result, many users subverted the system by creating single-valued "note" structures to store their information.

This is not to say that structure and formalism are to be avoided, indeed they could be said to be the goal of the analytic process. To address this, Shipman suggests designing systems that support *incremental formalism* [Shipman and Marshall 1999b]. Initially, users should not be asked to formalize - structure should be emergent with understanding. In other words, users should be able to create and destroy structures as their understanding of the underlying information changes.

Space provides a medium which can support incremental formalism. By positioning information in space, users can create structures based on perception rather than formalism. As Nakakoji says, "[the] perceptual external representations provide information that can be directly perceived and used without being interpreted and formulated explicitly" [Nakakoji et al. 2000]. For instance, two objects can be placed in proximity to one another to indicate some potential relationship without explicitly stating the nature of the relationship. As more

information is added, the nature of the relationship may become clearer. On the other hand, it may be decided that there is no relationship, in which case it is trivial to move one of the objects, thus breaking the perceived connection.

There are several examples of space being used explicitly to support incremental formalism. In the previously mentioned Aquanet example, users were observed arranging their notes spatially in these loose, dynamic structures that evolved with their understanding [Marshall and Rogers 1992]. The very basis of spatial hypertext evolved out of this concept [Shipman et al. 1995, Shipman and Marshall 1999a, Shipman et al. 2001a]. Bauer et al. report that users of Dynapad, their zoomable document organization system, developed many complex organizational schemes that "emerged gradually through reflective interaction with the collection" and were never fully articulated [Bauer et al. 2005]. Robinson also observed this in his studies. Analysts were freely mixing spatial metaphors in the space provided, combining lists, timelines, clusters and other structures together in a shared space, relying on perception and context to delineate the structures [Robinson 2008a].

2.3 Spatial Tools

A number of projects have sought to use space as a primary design element, either by allowing user freedom within it or by explicitly encoding meaning into it. In this section, we touch on a number of related applications that demonstrate some of the ways that space has been integrated into applications. While the work involving virtual environments cited above provides useful evidence for the benefits of large displays for spatial applications, the emphasis of this work is on 2D applications, as such, 3D applications are not addressed in this section.

2.3.1 Virtual workspaces

As we have repeatedly stated, humans are well adapted to interacting with a spatial environment. However, conventional displays provide fairly minimal interactive space. To work around this problem, researchers have developed a number of techniques providing more "space" than is available on the display [Furnas and Bederson 1995, Cockburn et al. 2009]. Henderson and Card summarized the various options as alternating screen usage (swapping between different visible sets of objects), distorted views (e.g., fisheye views [Furnas 1986]), large virtual workspaces, and multiple virtual workspaces [D. Austin Henderson and Card 1986]. The virtual workspace provides a large information space, and the user is provided a viewport that can be moved around the space (i.e., virtual navigation). We focus on the virtual workspace for a couple of reasons. First, it is a very familiar technique for providing additional space found in any image viewer, online mapping applications, and even word processing applications. Second, unlike the other techniques for providing more space, it maps directly to our use of the large display. The large display allow us to grow the viewport, and thus display more of the space. Of course, from a research perspective, our focus is on the point where the viewport is big enough to display the entire contents of the space, thus transitioning it from a virtual workspace into physical one.

While the members of this class are legion, we would like to highlight a couple of tools that have goals closely aligned to ours. The Pad++ system [Bederson et al. 1994] is a compelling example of this class of application. The goal of this project is to leverage users' spatial and geographic abilities to support navigation through a multimedia information space. Unlike most virtual workspaces, Pad++ attempts to offer the user a workspace that is conceptually infinite in resolution as well as extent, and to provide the experience of working with a single unified workspace through an emphasis on smooth panning and zooming. The

previously mentioned spatial document manager Dynapad [Bauer et al. 2005] was built on top of Pad++. The focus of this work is on spatial arrangement and quick visual access to information. In their observation of users working with prototype spaces, they found that several typical actions stood out: sorting, separating, adjoining and aligning. All of these involved the creation of *piles* (used here to indicate a visually distinct collection, not necessarily implying overlap). Dynapad supports this use of space by making piles explicit and through a number of spatial tools that can perform operations like sorting, clustering, and aligning.

While the zoomable interface creates a large virtual workspace, representation is still somewhat problematic in Dynapad due to the limited size of the viewport into the workspace. Images provide a natural thumbnail that can be zoomed to a full resolution representation, but other document types are more difficult. For papers, "enriched thumbnails", an integration of visually distinctive attributes of the document into a single page, were used, but no technique was provided to actually view all of the contents of a non-pictorial document.

Spatial hypertext is a special class of virtual workspace that is particularly close to our proposed work. The roots of spatial hypertext can be found in the early notes based systems like NoteCard [Halasz et al. 1987, Halasz 1988] and Aquanet [Marshall et al. 1991]. As discussed above, the primary problem with these tools was that they imposed too much structure on the space [Marshall and Rogers 1992, Shipman et al. 1995, Shipman and Marshall 1999a]. The tools were intended to be used for thinking and planning, but the requirement to develop structure before the problem space was even understood led to early ossification or subversion of the system.

Spatial hypertext was developed out of the observation that users were creating relationships and structures implicitly in the workspace by using spatial proximity and ordering. Thus,

spatial hypertext systems are, in essence, open workspaces that allow users to place colored text boxes and other objects into the space in ways that are meaningful to them. While most hypertext systems are really non-linear authoring systems, spatial hypertext is intended to be used for sensemaking, allowing the user to work directly with information, exploring arrangements in space as a way to construct meaning [Francisco-Revilla and Shipman 2005, Shipman et al. 2002]. There are now a number of tools based on spatial hypertext, including Visual Knowledge Builder (VKB) [Shipman et al. 2001a] and the commercial Tinderbox [Bernstein 2003].

While the proposed work draws heavily upon the theories of spatial hypertext, the research agenda diverges in a key way. The primary emphasis of much of the ongoing research in the spatial hypertext community is focused on the development of spatial parsers, which are designed to recognize spatial structures [Francisco-Revilla and Shipman 2005, Shipman et al. 1995, Yamamoto et al. 2005]. The goal of this work is to produce tools that can intelligently assist the user as structures become apparent by suggesting ways to structure new information or to find meaningful ways to "export" the contents of a spatial hypertext document to a different format. The technology that results from this research may prove useful for future sensemaking tools, but the work proposed here is more concerned with further studying how space can be used by analysts.

Virtual workspaces provide several advantages. The size of the space is bounded only by computer memory and performance constraints, rather than on the availability of physical space and display hardware. Moving a mile across a virtual workspace requires no more effort than moving an inch. Virtual workspaces can also be warped and bent to bring discontinuous regions together. However, the larger and more complex these spaces become, the more burden is placed on the user in terms of comprehending the space. Techniques, such as overview + detail and focus+context [Baudisch et al. 2002], have been developed to help

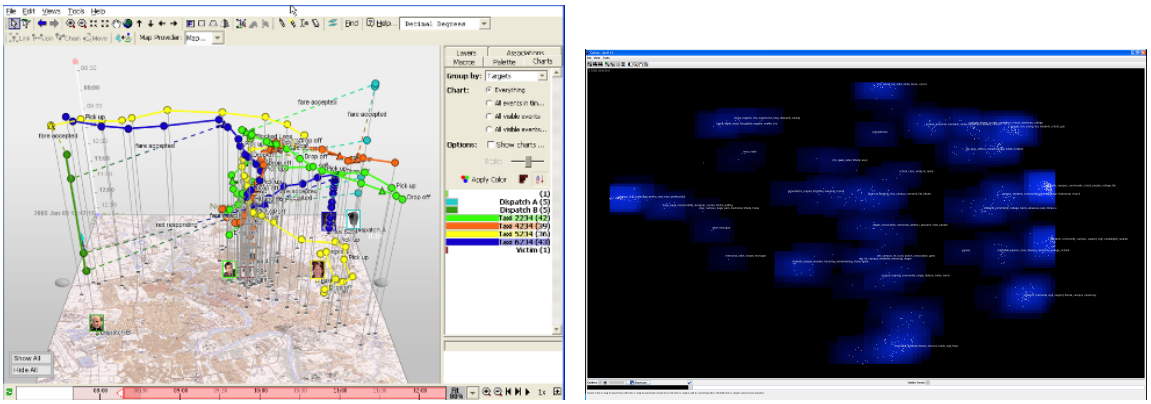

(a) GeoTime **(b)** IN-SPIRE

Figure 2.2: (a)GeoTime uses space to encode meaning such as geospatial information with overlaid temporal data in the third dimension [from [Eccles et al. 2007]] (b) IN-SPIRE's GalaxyView represents documents as points in space

users with this problem, but there is still cognitive effort involved in maintaining the mapping between the space and the current view. We are interested in looking at what happens when the space is human scale and provides its own model. In such a "real" environment, we feel that the abilities humans already posses for understanding and interacting with space will be more relevant and reduce the cost of working with the space, thus making it a better cognitive ally.

2.3.2 Sensemaking tools

There are a number of tools that are specific to the domain of intelligence analysis that make use of space in interesting ways. One common approach is for the system to impose the meaning onto the space. An example of this would be GeoTime, in which structured information is placed in a well-defined, three dimensional space. The primary use of the three dimensions is for events to be simultaneously mapped both geospatially (on a map) and temporally (in the space "over" the map) [Kapler and Wright 2005] (Figure 2.2a). While originally designed for geospatial information, GeoTime does also allow less constrained arrangements on the base plane, such as arbitrary graphs. A slightly less structured example

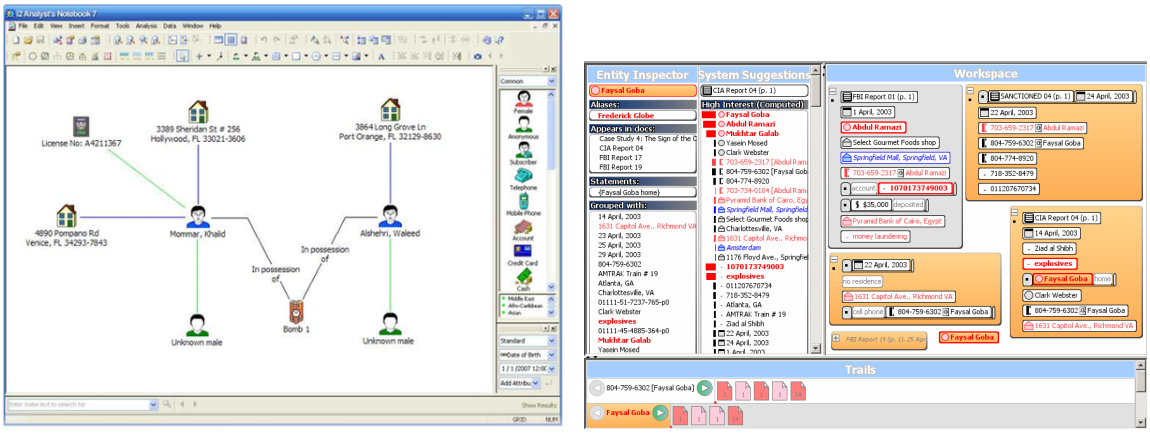

(a) Analyst's Notebook **(b)** Entity Workspace

Figure 2.3: (a) An example link analysis graph from Analyst's notebook [from [Ana 2011]] (b) Attributes are collected together to create entities in Entity Workspace [from [Bier et al. 2006]]

would be IN-SPIRE, which plots documents to points in space based on document similarity (Figure 2.2b). The user has no control over placement, but any meaning associated with a region of space is assigned by the user based on the perceived similarity of documents present in that location [Wise et al. 1995].

An increasingly common approach is to visualize intelligence data and products as node-link graphs, such as those found in Analyst's Notebook [Ana 2011] (Figure 2.3a), Sentinel Visualizer [sen 2011], and Palantir Government [Pal 2011]. These tools use the space to construct graphs that connect together people, places and organizations based on automatic document parsing or guidance provided by the analyst. A potential problem with this approach is that over-simplification can be a problem since most of the details and context are hidden, making these potentially more appropriate as report tools than thinking tools [Wright et al. 2006].

The Entity Workspace [Bier et al. 2006; 2008] is a different kind of entity-based system that provides an open workspace that can be arranged by the analyst. The difference is that all of the information is broken down into highly structured entity objects. These entity objects can

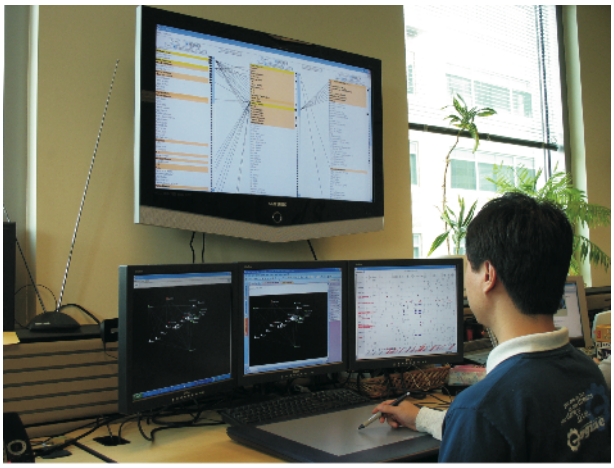

Figure 2.4: Multiple displays help the analyst to keep track of Jigsaw's multiple views [from [Stasko et al. 2007]]

be dragged around and placed anywhere. In addition, the objects can be "snapped together" to create composite entity statements that indicate relationships (Figure 2.3b). While this tool has the potential to use spatial relationships to express meaning, the spatial constraints on the workspace due to the limited display space available on conventional systems mean that there appears to be little room for many relationships to be expressed and the emphasis of their research is not currently focused on supporting that type of exploration.

Another approach to an entity-based workflow is provided by Jigsaw [Stasko et al. 2007]. The primary goal of Jigsaw is to facilitate the exploration of relationships between entities while maintaining an analyst-driven workflow. The tool provides the analyst with a collection of linked views that present information about entities at different levels, such as which entities are connected in some way, which documents an entity can be found in, as well as a view that shows the actual reports with entities highlighted. One interesting feature of Jigsaw is that significant views can be duplicated and frozen, leading to a proliferation of windows. Stasko et al. report that screen real estate is a primary issue and that ideally, to reduce excessive window management, a multi-monitor environment should be used (Figure 2.4).

Sandbox, the sensemaking half of the nSpace suite [Wright et al. 2006], is a variation of the

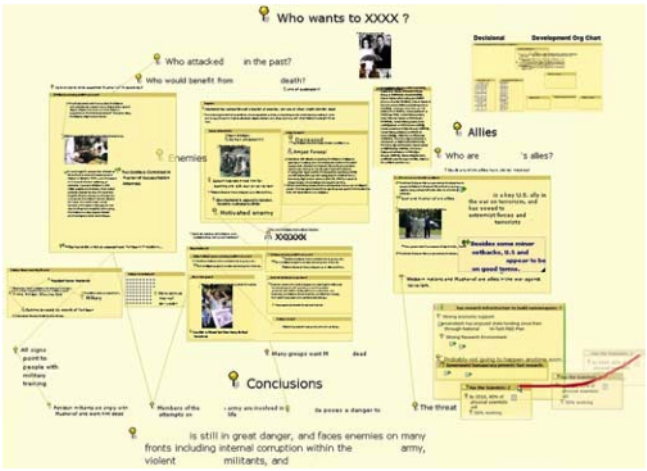

Figure 2.5: Sandbox's freeform workspace allows text and document snippets to be freely arranged in space. [from [Wright et al. 2006]]

spatial hypertext approach with the addition of integration with analytic tools. Sandbox provides the analyst with a flexible, open workspace where he or she can freely move data around, organizing, annotating, hypothesizing, etc. The space provides some specialized tools for activities like the analysis of competing hypotheses, but it primarily a free-form space that requires the analyst to provide meaning (Figure 2.5). In many ways, Sandbox was developed to leverage the same benefits of space that we have identified. However, the tool is targeted at conventional displays, making the space it works with almost entirely virtual, so it lacks some of the advantages of working with real space, such as visual and spatial persistence, and the bias towards egocentric frames of reference. There is no reason that it could not be used on a large, high-resolution display, but the use of the virtual workspace metaphor has determined a number of features of the environment, such as the use of short snippets of text, reduced size representations, and the use of the non-spatial half of the nSpace suite for foraging. All of these features could be either unnecessary or potentially constraining on a large, high-resolution display.

2.4 Summary

In the course of our initial exploratory studies of large, high-resolution displays and the sensemaking process, we determined that the physical attributes of the display created an implicitly spatial environment. Analyzing how that space was used led us to identify memory and organization as the two major forms of support provided by the display for the sensemaking process. In this chapter, we have returned to the literature to find support for these observations as well as potential explanations. We showed how the perception of space is created through the physical attributes of the display, as the scale biases the user into adopting an egocentric frame of reference and the large quantity of available pixels encourages the use of physical navigation.

In examining the role of the space for sensemaking, we have illustrated how the physical properties of the display combine with the user's embodied resources to facilitate memory – making use of aspects of the perceptual system, spatial awareness and spatial memory to provide external memory. The semantic advantages of space are rooted in the human ability to easily perceive spatial patterns and interpret them as relationships, and the flexible, freeform nature of space, which permits the user to externalize arbitrary and even unspecified relationships into it. We have also stressed the importance of incremental formalism, its role in supporting the development of understanding, and the way in which space supports this technique.

In the next chapter, we will discuss the original studies that revealed these key areas of exploration.

Chapter 3

The Role of Large, High-Resolution Displays for Sensemaking

Our theories about the spatial affordances of large, high-resolution displays and the role of the available space in the sensemaking process are rooted in a pair of observational studies of users solving analytic problems. In some ways, these studies closely resemble Robinson's examination of analysts performing synthesis using physical artifacts [Robinson 2008a]. Robinson's study, however, focused primarily on how analysts used space to synthesize pre-selected data and how space was used collaboratively to negotiate a shared meaning. Our studies were situated on a large, high-resolution display, allowing us to also examine the impact of technology and broaden the focus to the whole sensemaking process.

The goals of these studies were twofold. The first goal was to determine if the size of the display did foster a spatially based approach. While the tools we provided were simple windowed text editors, our anticipation was that the relative size of the display would change the

subject's perceptions so that they treated the space as a free-form organizational environment similar to that explicitly provided by sensemaking tools such as Sandbox [Wright et al. 2006] or Visual Knowledge Builder (VKB) [Shipman et al. 2001b]. The second goal was to examine the actual mechanics of sensemaking in this environment. In other words, we were interested in how users would make use of the space at the various stages of sensemaking, and how it affected the overall sensemaking process.

The Environment

For both studies, the large, high-resolution display was running Windows XP. No special analytic tools were provided to our participants. Instead, the study participants only used basic text editors (e.g., WordPad), image viewers (e.g., Windows Picture and Fax Viewer), and Windows Explorer for file browsing and search. While this presented some challenges to the participants, it allowed us to look specifically at the properties of the display itself.

Data set

The analytic problem presented to the study participants was a data set originally developed for the interactive session of the VAST 2006 contest (code name "Stegosaurus") [Grinstein et al. 2006]. This data set consists of:

- 230 simulated news stories
- 3 images
- 1 map
- 1 spreadsheet

- 3 reference documents (background information)
- 1 XML file containing preprocessed data from the data set

The preprocessed data was produced by performing entity extraction on the text documents in the dataset. The file contains entries for various people, locations and organizations that appear in the dataset and provides a list of the documents in which they appear. In addition to these materials, we created a second preprocessed version of the news stories that consolidated all 230 of them into a single document. The primary goal of this file was to allow subjects to use the somewhat more robust search tools available within text editors rather than having them rely entirely on the search tool available in Windows Explorer.

Only about ten documents out of the 238 are actually required to establish ground truth. Several more provide background information that is not necessarily required to solve the scenario, and the remainder has nothing to offer the analyst but distractions. The collection also contains a significant red herring, which increases the analytic complexity of finding the underlying plot. Another feature of this data set is a 'seed' document. The scenario begins with a briefing from the sheriff of the fictional town of Alderwood asking for help regarding a strange incident in town, providing a clear start point.

There were several benefits to using this dataset for our exploration. It was originally designed so that analysts could be expected to make reasonable progress, if not solve the scenario, at the VAST conference within a two hour session. The seed document also plays a significant role in this as it focuses the investigation rather than requiring the analyst to consider the data set as a whole. Despite this, the scenario embedded in the data set is complex and requires a fair amount of intuition to piece the various clues together. There is also a known ground truth, which makes it easier to evaluate progress and the depth of insights arrived at by study participants.

Methodology

While we did examine the analytic results produced by our subjects, our primary interest was in the process. Taking the perspective of distributed cognition, the real focus of our analysis was the generation and transformation of observable representations that represented the communication between the analyst and the computer. Where possible we drew connections back to Pirolli and Card's sensemaking loop [Pirolli and Card 2005] (Figure 1.1).

Each subject worked alone, and our analysis is based on a mixture of screen captures taken at regular intervals throughout each session, video, and observations made by the researchers. We specifically did not ask our subjects to use a 'think aloud' protocol during the sessions to prevent cognitive interference with the sensemaking process [Boren and Ramey 2000]. Following each session, we carried out semi-structured interviews with the subjects. In these interviews, our main questions were: What conclusions were reached? How were they reached? How did they use the space? What were their general reactions to working in the provided environment?

3.1 Display Size Comparison

This study was designed to explore the effect of display size on the sensemaking process and display usage. We recruited eight students to act as analysts and work through the VAST scenario. Half of the subjects were provided with a small display (a conventional 17" display), while the second half used the large display. For convenience, we will label the subjects based on the condition they were in, so subjects S1-S4 used the small display, while subjects L1-L4 used the large display. Of the eight subjects (7 male, 1 female), four

of the participants had some exposure large displays, but none of them used them for daily work and only two participants typically used more than a single monitor. In addition to the displays, participants were also provided with blank paper, writing implements and a whiteboard to use if they thought that it would be useful.

3.1.1 Results

Quantitative results

Two quantitative metrics were collected during this study, time and correctness. While each session was timed, the participants were told that their goal was to solve the scenario and to be able to back up their conclusions with evidence. To score the results of this study, we used a modified version of the metric proposed by Plaisant et al. [2008]. The metric breaks down the results into three separate scores for who, where and what/when. There are four important metrics: true positives (T_p), true negatives (T_n), false positives (F_p), and false negatives (F_n). Using these, the three scores were calculated as:

$$\begin{aligned} who &= T_p + T_n - \frac{1}{2}(F_p + F_n) + .25 \\ where &= T_p + T_n - \frac{1}{2}(F_n) \\ what/when &= T_p + T_n - \frac{1}{2}(F_n) \end{aligned}$$

These scores were then summed for the final accuracy score. Several subjects achieved perfect or near perfect scores, but, in general, the scores varied significantly. Times were similarly varied, with the fastest solution taking just under 1.5 hours and the longest solution taking

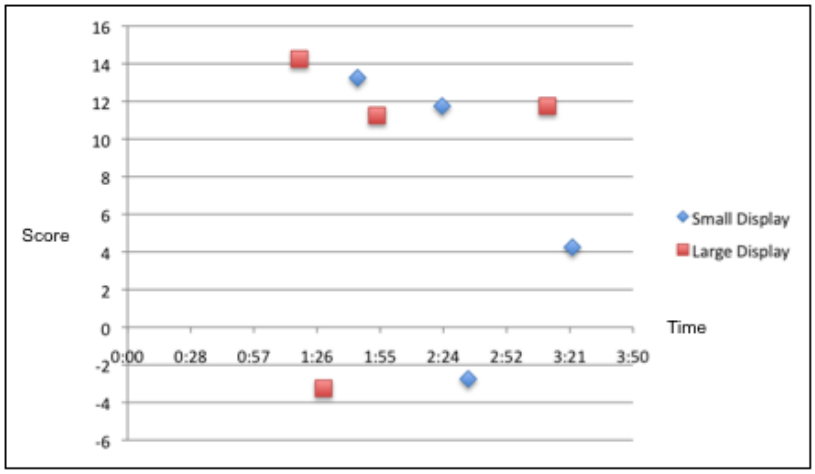

Figure 3.1: Graph of score plotted against time showing the wide variation

close to 3.5 hours (Figure 3.1). Given the relatively small subject pool, the varied backgrounds, and the complexity of the problem, it is unsurprising that there were no significant differences between the large and small display conditions

Observations

While there were no significant differences between the two groups with respect to time or completeness of the solution, there were a number of key behavioral differences between the two groups.

As one might expect, the primary difference lay in the approaches to document management. On the small display, almost universally the subjects maximized every document window and used the taskbar to switch between them. As the names of the documents were primarily numerical, the one representation available to this group was particularly unhelpful, so there was considerable thrashing (paging rapidly through open windows), every time one of the small display subjects tried to switch to a different document.

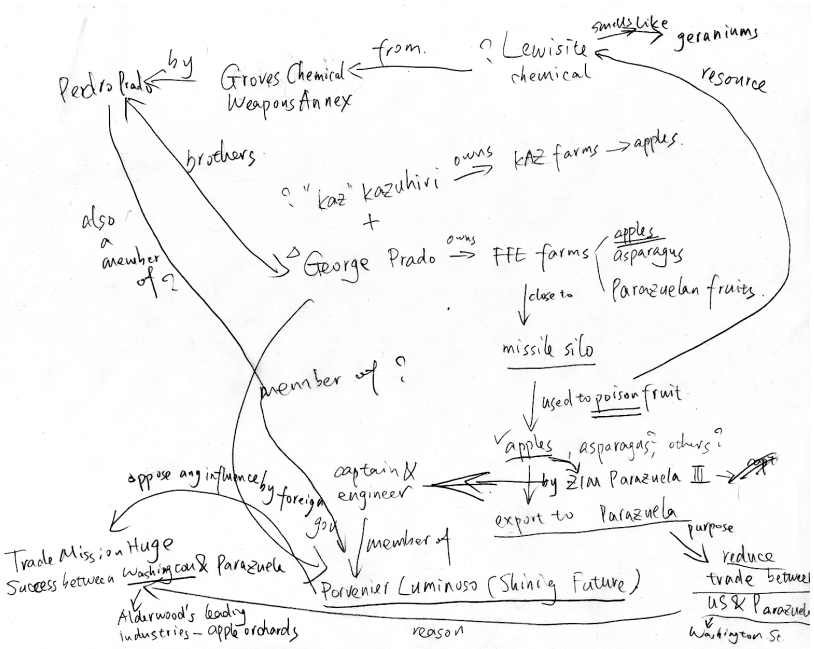

Figure 3.2: Subject S1's diagram of the key elements of the scenario

The large display group, on the other hand, used the available space to lay out their documents. Unfortunately, the degree to which they could use the display was limited by the approach that they took. With the exception of subject L4, all of the large display subjects used the consolidated document for both search and reading and never opened up the individual news articles. The use of the space then was restricted to displaying the background documents. However, it was clear that this was still an important use of the space. For example, subject L1 spent the first 45 minutes of the study laying out his space. Rather than directly addressing the problem, he started by reading all of the background documents, summarizing them if necessary to make them fit on the display so that they would all be available while he was pursuing the investigation.

Of more interest is the marked difference in the representations produced by these two groups. The small display group all produced two different classes of artifacts. The first were digital notes, kept in a new document in the file system. Predominantly, these contained whole articles copied out of the consolidated document, and seemed to serve as a quick

reference file so that the subject could rapidly revisit articles that seemed interesting. The second were physical notes, kept primarily on paper. These notes contained a variety of intermingled representations: keywords, names, questions they wanted to answer, leads to pursue, and theories that they were forming. We can map all of these representations to various stages in the sensemaking process. The digital files maintained by these subjects formed a rough shoebox – documents that they had identified as being important. The physical notes, on the other hand could be seen as a loose collection of evidence snippets and mnemonics for potential schemas. Many of these were fragmentary – single words or phrases, usually broken into sections based on elements of the investigation. In two instances, the subjects also produced rough diagrams that were essentially expanded social network graphs illustrating the relationships between people and organizations visually (Figure 3.2).

The large display group, on the other hand, with the exception of subject L4, produced no paper artifacts; however all of them produced digital notes. The notes taken by this group are marked by their sparseness, and they bore more in common with the paper notes taken by the small display group rather than the digital files created by that group. The information represented in these notes can be largely classified as being the evidence file of this group, consisting largely of short snippets and keywords. L4's use of paper can probably be attributed to lack of experience. She started by writing out notes on paper, but soon switched to taking notes on the display, keeping the two versions coordinated for some time before ultimately abandoning the paper.

The primary effect that we observed is the use of the space as a form of easily accessible external memory. Spreading the documents out in the space is a simple example of making information visibly available for consultation. Physical navigation (glancing around) is augmenting memory and replacing virtual navigation (flipping between document windows).

This is similar to the effect observed by Ball in his studies of geospatial visualization [Ball et al. 2007].

The various types of representations created by these groups can also be tied back to memory. The notes containing snippets and thoughts produced by both groups are clearly important. They represent the evidence and reasoning being followed by these subjects. The creation of these representations can be tied directly back to one of the "leverage points" identified by Pirolli and Card – the span of attention for evidence and hypotheses [Pirolli and Card 2005]. As an investigation proceeds, there is a need to somehow manage the growing collection of evidence and its relations. For both groups, it was a priority that these representations were available and editable while they were reading and researching, which explains why the small display group felt it necessary to produce these externally to the display due to the lack of space.

We can tie the digital representations produced by the small display group back to a specialized use of memory – memory for context. Both groups made extensive use of the consolidated document, but only the small display group copied whole articles out of the text into a shoebox file. A possible explanation for this can be found in how the consolidated document was displayed. On the small display, the document was maximized to fit the width of the window, limiting the visibility to about one half page of text. On the large display, however, most of the users spread the document window out and displayed the document two-up (two full pages simultaneously visible on the display). This created a marked difference in how the two groups moved around the document. The small display subjects frequently got lost while they were searching the large document. They did not realize either that there were more matches to be found, or that a search had wrapped around and started again at the beginning. As a result, the large display group was markedly more confident in their ability to remember where they were and to rapidly re-find articles that they had read previously.

3.1.2 Discussion

The main results from this study is the observation that the display size changed subject behavior. More specifically it altered the kinds of representations that were being produced and how they were manipulated. The first observation that we can make is that all of the subjects felt it necessary to keep some kind of record of significant names and keywords as they progressed through the scenario. It is important to note that while all subjects took some variation of these notes, the form differed between the two conditions. The key difference appears to be one of visibility. On the large display, an extra file containing notes could be persistently visible without interfering with other documents that are open and being read. On the small display, on the other hand, most users maximized documents to fill the display, so digital notes could only be visible when the source document was hidden. Physical notes, however, allowed the subjects to create persistent representations that could be used simultaneously with the source documents.

A further point to make here is the issue of malleability and representational power. By using only text files to store notes, the subjects using the large display lost the ability to create complex representations like social network graphs. On the other hand, in addition to the advantage of the notes and source documents sharing the same environment physically and technologically, the digital nature of the notes provided a level of malleability not present in physical representations. For example, if we examine Figure 3.2 again, we see that as the investigation progressed, space started to come at a premium and the subject began to cram entities in at the bottom of the page. We can also see that many entities were written onto the page before their relationship to the rest of the entities became obvious, leading to a tangle of relationships. Rearranging the diagram would help to clarify the conclusion, but at the cost of redrawing it from scratch.

One unexpected result of this study was that there was less use of the space on the large display for encoding relationships and meaning than we expected. Our anticipation was that we would observe subjects arranging documents spatially to encode meaning into the space, but we saw very little of that behavior. Most of the arranging was done either to optimize workflow or to optimize visibility without much concern for relationships. This can largely be attributed to the inclusion of the consolidated document in the dataset. Both groups made extensive use of this document for searching and reading, with very few subjects ever opening the individual documents. Subjects appeared to be more comfortable using the search facilities of Word rather than the document contents search provided by Windows Explorer, and once a document was found, it was easier to read it in place than to search for it again in the file system, so few users left the consolidated document. Since there was no other reason to open individual articles, there was no opportunity to use them as representations.

Given this behavior, it is interesting that the large display group never copied whole articles out of the consolidated document. It is possible that since they could see more of the document (most large display subjects opened the document so that they could see two to three whole pages as opposed to the 2/3 of a page available on the small display), the large display subjects felt more comfortable that they could rapidly re-find important articles.

Another observation that we made during this study was that the reference documents stood out from the rest of the collection (they are named differently, in a different directory, and are .doc format rather than the plain text used by the news articles). This difference led some subjects to read them first rather than "discovering" them at appropriate points in the investigation through search, which, because of their specific nature, biased the investigation towards chemical weapons and South American terrorist groups.

While this study did not provided clear evidence of some quantifiable advantage to using the large display for sensemaking, it did illustrate how the different displays change behavior. The study also suggests that to encourage better use of the available space, users need to be provided with better representational support.

3.2 Analyst Study

For the second study, we had the opportunity to recruit five subjects from a government laboratory (A1-A5). Four of the participants were practicing analysts from a variety of fields, and the fifth was a developer of analytic tools. Unlike the first study, this was not a comparative analysis. Considering the results of the initial study and the limited available subject pool, it was decided that more would be learned by closely observing the use of space than trying to compare performance across the two systems.

The primary goal of this study was to observe domain experts in this environment. As Kirsh points out, it is experts who develop the most sophisticated spatial strategies [Kirsh 1995]. While the recruited subjects were not experts in the environment (i.e., the human scale display), they were domain experts with foreknowledge about how to approach analytic problems and pre-established techniques for organization and identification of objects of importance. Our interest was in examining how these techniques translated to the environment.

The basic methodology of this study was the same as the one used for the prior comparative study, but with a couple of key differences. The primary difference was that all of the subjects used the large display, and we did not score the solutions since we were more interested in the process than the solution. In addition, the sessions were shorter due to

the time constraints of the subjects and each session included an informal interview about general work practices, experiences completing the scenario, and desired features for future spatially oriented sensemaking tools.

To address some of the problems with the earlier study, we also reorganized the dataset. We removed the consolidated news document and the XML document to encourage an increased number of open documents, providing more opportunities for spatial strategies to develop. The background documents were also converted to plain text and integrated into a shared directory with the news articles to prevent the analysts from focusing on them prematurely.

3.2.1 Results

Overall Reaction

Reaction to the display was overwhelmingly positive. One subject told us that she initially thought the setup was "way over the top", and doubted she would find a use for more than one of the monitors, but was totally converted by the end of the session: "virtual organization really rocks!"

Initially most of the subjects had similar reactions– they were thinking in terms of opening and closing windows and could not imagine how they would use all of the space. However, in a relatively short period, the study participants generally came to regard the environment as something distinctly different from just a larger display or extra monitors, despite the use of conventional Windows tools. One analyst commented that the environment was less like a computer screen than an environment for video games. Another compared the space to the conference table he typically uses to lay out papers for sensemaking. Almost universally,

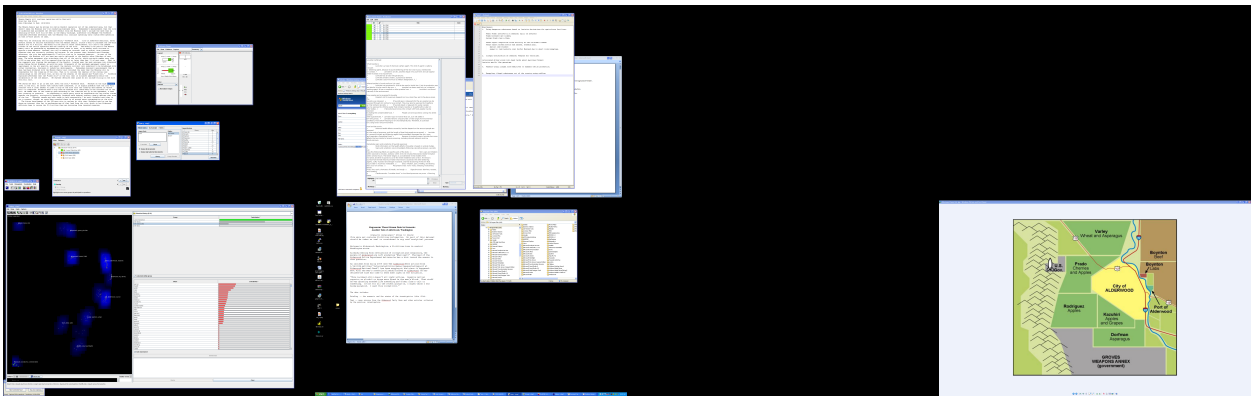

Figure 3.3: Subject A5's workspace showing how the use of IN-SPIRE altered the use of the space

the analysts changed their use of the available tools to leverage the spatial nature of the environment. This shift was primarily marked by a change from regarding windows as document viewers to actual documents that could be spread out over the display and moved around as part of the sensemaking process. These reactions agree with our intuition about how display size and resolution change user perceptions.

Observations

With one notable exception, the more spatial environment had a clear impact on how the subjects approached the problem. The exception was not an analyst, but a developer of analytic tools. After performing one search with Windows Explorer, he discovered a copy of IN-SPIRE [Wise et al. 1995] on the experimental machine and proceeded to use it for the rest of the session. The rational for using IN-SPIRE was the improved contextual information, clustering, and superior search facilities.

Non-spatial use of the display IN-SPIRE is a specialized intelligence analysis application for analyzing large document collections, and provides tools for document clustering, searching and viewing with embedded contextual information. The use of IN-SPIRE changed

the space from one dominated by information (many open documents) to one dominated by tools (multiple views and panels), which changed the dynamics of its use. The most obvious evidence of this is the rather limited use of the available space (Figure 3.3). Most of the work has been constrained to half of the display. The only item shown on the second half of the display is the map, which has been blown up from its normal dimensions of 542×515 to fill one entire monitor. This behavior harkens back to the impact of the consolidated document used in the earlier study.

An important point to make is that the use of space to fit multiple tools and the use for document organization is not an either/or proposition. In this situation, IN-SPIRE was chosen for its superior search capabilities. As it happens, the integrated query tool / document viewer only permits the user to view the contents of a single document at a time. The subject could have then hunted down the original file and opened it separately, but that would have required a fair amount of extra work. It is telling that this is the only subject who maintained a separate notes file to keep track of his thoughts and progress through the scenario. The lesson here is that these spatial environments create new possibilities that were not anticipated and are frequently not well supported by tools designed for conventional displays.

For the remainder of this section, we shall discuss the rest of the subjects who did use the display spatially and how it affected their approach to the problem.

External Memory As with the large display group from the comparative study, a primary use for the space was as a form of external memory. All of the analysts followed the same general pattern as they searched. They would enter a keyword into Windows Explorer and then scan through the results using the built in preview function to skim each file in

turn. If the document seemed interesting or relevant, the document was opened. This exactly mirrors the initial stages in the sensemaking loop, with the display assuming the role of the shoebox. Frequently, the importance attached to the document by the analyst could be judged by how the document was placed into the space. Important documents tended to be tiled, with their contents fully visible. Documents with potential, but no clear importance were piled. The result of this was a partitioning of the document space into three classifications: "important" (fully visible), "interesting" (piled, easily relocated), and "unused" (not visible/opened, harder to relocate). There was, of course, an implicit fourth classification, "seen and rejected", that was indistinguishable from "unused". This did cause some problems for some of the subjects as they "found" several documents multiple times through different searches and had to spend some time reexamining them to make sure that they had really seen them. None of the subjects came up with a strategy that addressed this problem.

We observed several behaviors that supported the interpretation of the space as external memory. One behavioral indicator was the glance while reading. While reading one document, the subject would turn briefly to consult a different document, usually to check a name, or to consult the map. This is clear evidence of a readily available form of external memory.

Another key behavior was the rereading or rescanning of the visible information. For all of our subjects, this behavior was manifested during pauses in the investigation. These pauses happened throughout each session, but with increasing frequency towards the end of the session. Each analyst would just scan back over the documents, sometimes rereading, and sometimes just skimming the document. It is likely that this process is part of the schematizing process, connecting together some of the disparate pieces of information into

potential scenarios. This behavior, of course, is made possible by the visual availability of the documents at a high level of detail.

One particularly interesting moment came out of one of these periods of rescanning. One of the analysts was scanning back over some older documents and suddenly recognized a name that had been mentioned in a more recently read document. The name had not seemed important previously, so it was quickly forgotten, but rediscovering it forged a link between the two documents and encouraged the analyst to (correctly) focus on this person as a "person-of-interest." This serendipitous event happened principally because the documents were visible and available at a high level of detail. While an alternative, lower fidelity representation may have been enough to remind the analyst of the overall contents of the document, it would not have been able to provide her this link, which was forged on what was initially a seemingly unimportant piece of information in the document.

It is also interesting that the subjects created very few other external representations. Only the participant who used IN-SPIRE took notes, and only one other participant created a new text document. In this case, she was explicitly writing out a timeline of events rather than laying out the documents. The only additional representation of interest was the creation of a second search window. In this instance, the subject wanted to pursue an interesting line of investigation without losing the current set of search results (in essence, the investigation was branching). This could have been accomplished by writing down the search term, but it would have lost much of the contextual information (the look of the window, the highlight indicating the last document that had been explored, and the number of documents that were still unexplored).

The implication of the relative lack of explicit memory aids is that the presence of the document contents and the rough categorizations was enough to keep track of the investigation.

This is also evidenced by the behavior displayed by the subjects during the session. As mentioned previously, the subjects frequently took moments to pause and scan over the contents of the display. This is clearly an act of refreshing, revisiting the documents and the organization to see if a coherent picture was emerging. While less detailed representations may have served in some of these cases, the presence of the actual text was important. In one notable instance, one of the subjects was scanning over some previously read documents and suddenly noticed a name that had come up repeatedly in the investigation. The name provided an important link between the old document and the current line of the investigation, but it had not been noticed or remembered on the first reading because the name was not yet important.

Atomizing Information Returning to the sensemaking loop, we note the importance of information *atomization* or *extraction*. This is the process of identifying and isolating pieces of evidence from documents in the shoebox to create an *evidence file*.

None of the subjects explicitly produced a separate evidence file. The closest direct analog was the notes produced by subjects L1-L5 in the comparative study, and by subject A1 in this study (none of the other analysts created a separate notes file). The atomized information is manifested in these documents as names, keywords and phrases taken from the documents.

This is not to say that the analysts did not atomize the data. Rather than extracting it, they all isolated it within the documents through highlighting. This was clearly an important activity because all of the analysts did this, despite the difficulties it entailed. The problem is that proper highlighting is not available in WordPad. To work around this, the analysts tried a number of alternative techniques, such as changing the color of the text or bolding it to make it stand out. Most of the analysts even discovered that they could make semi-persistent highlights just by selecting some text and then not touching the document again

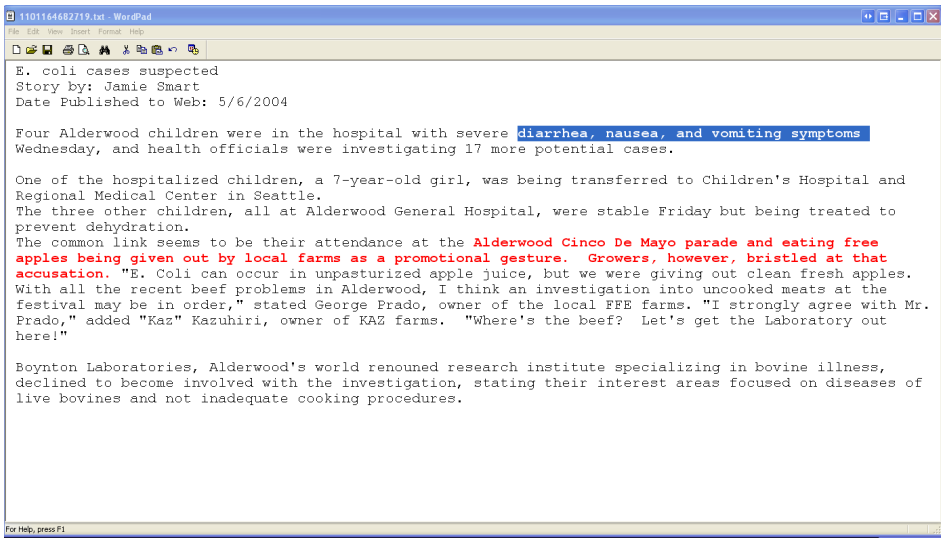

Figure 3.4: Document showing both selection highlighting and text color highlighting

(Figure 3.4). All of these workarounds suggest just how important they found these visual representations.

Organizational Strategies One common thread that ran through the behavior of all of the subjects in both this and the pilot study was the establishment of a 'work zone' - a region of the display that served as the primary focus. This was typically in the middle of the display and frequently towards the bottom. This is where the active search tool resided and where much of the reading was done. The rest of the display was largely used peripherally, available for organization and consultation. This is largely explained by the fixed nature of the search tool. There was no reason for the subjects to move the search window around, so it anchored the primary activity (searching) to a fixed location on the display.

Once the analysts began to treat document windows as documents, they began to develop different organizational strategies. We can regard the most organizations and the resulting structures as a form of *evidence marshaling* and *schematizing*. The exceptions were small, localized organizations that arose as artifacts of the process – linear orderings that reflected

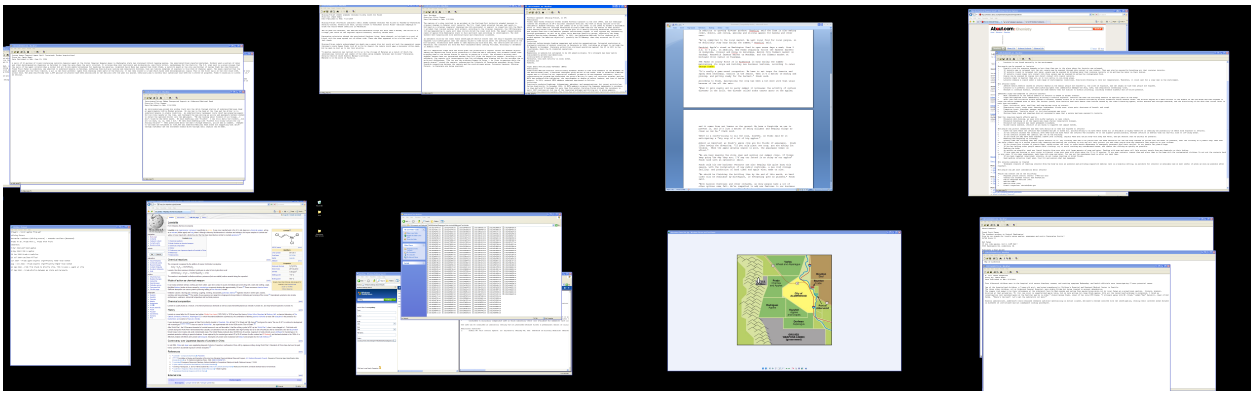

Figure 3.5: Final workspace for subject A1. Top row: 1) critical docs, 2) Shining Future, 3) George Prado, 4) chemical weapons. Bottom row: 5)timeline and wiki page, 6) work zone, 7) map, 8) potentials

the order in which documents were found. While not part of the marshaling process, this still provides valuable context with potential insight into the thought process that led to the discovery of the documents.

Most of the organization tended to take the form of clustering, or rough categorization (Figure 3.5). Interestingly, these categories ranged greatly in specificity. Some were quite general (e.g., 'potentially interesting' or 'background'), while some were quite specific (e.g., documents that discuss a particular individual or organization). Occasionally, we observed these categories evolve over time as more information was acquired. For example, subject A1 started a cluster that she referred to as 'background' information. However, as she learned more and placed new documents in the cluster, the meaning shifted to 'critical documents'. The contents did not change - her interpretation of them shifted. An interesting characteristic of the clusters created by the analysts is that some were carefully arranged to maintain visibility, while others (like the 'potentially interesting' cluster) were typically allowed to accumulate in a pile, with just enough left showing to trigger recognition. This clustering behavior is particularly interesting because it is a common technique for making sense of physical, rather than digital, documents [Russell et al. 2006]. When asked, the analysts commented that being able to organize documents directly in the space was replacing their

previous techniques, such as filing documents individually into directories, and printing them out and organizing the papers physically, either spread out on a table or into piles on the desk, showing how the environment has changed the analyst's perception of the documents and the workspace.

There are a couple of reasons why clustering formed the primary spatial metaphor. Categorization is one of the most basic forms of marshaling. This form of data immersion is a common approach to sensemaking when nothing is known about the situation, making it harder to apply precedent or more reasoned approaches [Heuer 1999]. It is also low cost and relatively easy. It may also be that the nature of the scenario did not require much beyond the identification of the important documents. While the documents are dated, understanding the temporal sequence is not necessary to make the required connections. It is somewhat important to understand the geography of the region, but the included map makes that fairly straightforward, and the geographic connections are simple enough that there is little need to explicitly link the documents to the map. In addition, there is little opportunity in the solving of the scenario to form alternate hypothesis and distribute evidence between them.

The structural properties of the display must also be acknowledged as playing a role in this direction. Monitor boundaries create natural partitions Grudin [2001]. While the tiling of the 'analyst's workstation' and the use of 30" panels minimizes the effects somewhat, the bezels that form the monitor boundaries are still present. In addition, there is relatively minimal support in Windows XP for window manipulation. On conventional displays, as mentioned previously, the normal document window is very close in size to the size of the containing space. So it is unsurprising that the tools are limited with respect to working with space. Looking back at Figure 3.5, we find a large amount of "wasted" space caused by documents that are consuming more space than they require. This is largely caused by tools that use history to determine the size of new documents rather than the contents, which

consequently resulted in a large number of short documents opening in sparsely populated windows. While these windows are easily resized, it is a step that few subjects bothered to take.In addition, in a spatial environment, the inclination is to perform operations on multiple objects simultaneously. For example, using physical sheets of paper, a single gesture is frequently all that is needed to stack a collection, spread out a pile, move an entire set around, and even sweep a region clear. As none of these operations are available, it is difficult to create and maintain anything beyond the simplest of structures. The one timeline that was constructed was actually constructed backwards (time advancing from right to left) because the analyst noticed that the first three documents that made up the structure happened to aligned that way in the space already, so to was easier to just place the forth one on the left than to reverse them all to flow the 'correct' direction. Updating the structures with new information becomes increasingly difficult in this environment (e.g., trying to insert a document into the middle of a sizable timeline).

Despite all of these factors, clustering was not the only organizational structure that was observed. One interesting structure was the aforementioned timeline created by subject A4 that flowed horizontally across two thirds of the display, crossing two bezels. All of the news articles were dated, and this analyst used the dates to order interesting articles as she uncovered them. Structurally, the timeline was somewhat rough. The documents mostly lined up, but not always, and some documents were allowed to overlap others. Seemingly, the most important rule was that the left edges of all of the documents were chronologically ordered.

The most interesting aspect of the timeline is what the analyst did with it towards the end of her investigation. The timeline was formed in strictly chronological order with no regard for content. In order to make sense of this collection of documents, the analyst pulled apart the timeline, categorizing the documents based on theme. This action did not break

Figure 3.6: Subject A4's timeline before (top) and after (bottom) restructuring to encode categorizations within the structure

the timeline, however, because the categorization happened along the vertical axis, with temporal ordering being maintained horizontally. In other words, the structure became an amalgamation of two simultaneous spatial metaphors (Figure 3.6).

Incremental Formalism The evolution of the clustered timeline is a clear example of incremental formalism, and also served to illustrate some of the difficulties inherent in trying to automate this process. Initially, the structure was a loose cluster of aligned documents. As more documents were added to the cluster, the analyst decided that she wanted to organize them chronologically. Coincidentally, the documents were already in order, albeit in reverse chronologically. As described above, rather than reorder them, she continued to build the timeline backwards. An important point, however, was that for an outside observer the structure did not change – but the analyst's interpretation of it did and it affected her understanding of the documents and how she interacted with that space from that point onward. While the other subjects did not have as clearly identifiable transformation of structure, only one subject explicitly considered how the space would be used at the beginning of the session. For most, the process was evolutionary. The typical procedure followed by most subjects was to browse through search results using the preview pane built into Windows Explorer. When documents that seemed interesting were discovered, they

were opened and "filed" in the space. This is a rough process, driven initially by the excess of free space. However, as more documents were placed in the space, regions begin to have meaning. The early regions were very general - categories are formed like "background", or "potentially interesting". As the investigation proceeded, more specific categories emerged, such as articles associated with a particular person.

This behavior is clear evidence of the use of the space for incremental formalism. While no strict structures emerged and some of the boundaries between bins were fairly loose (when not constrained by the bezels), at the end of the session, each subject could describe the organization of the space at a level that was abstracted away from individual documents.

Integration While we are able to pinpoint specific instances of the various stages of sensemaking throughout the sessions, we also note the high degree of integration that can be observed throughout the process. We observed both integration of process and integration of representation.

Integration of process reflects the fluidity of the process. The analysts we studied moved freely around the sensemaking loop, jumping through various levels of abstraction. For example, it was common for analysts to mix together skimming, extracting, and marshaling. New searches could be inspired by a phrase found in a new document or by a desire to find evidence to support a hypothesis. While interesting to observe, this form of integration is implied by Pirolli and Card's model.

Just as the various stages of the sensemaking process fluidly combine, so did the various representations. The use of highlighting is a prime example of this. Highlighting passages in a document is a form of identification and extraction. Many forms of atomization completely separate the snippet from the document (e.g., copying the passage into a new notes

document). Highlighting has the benefit that it isolates without removing the information from context. Highlights serve a second purpose by creating a richer representation for the document as a whole as well. They provide a visual cue that aides recognition of the document. As one analyst remarked, he "just need[ed] the pattern of the highlights" to recognize a document.

The documents themselves are also examples of integrated representations. At one level, they are containers of detailed information: the actual text and highlights. They can also be regarded holistically as representational proxies for their contents like an icon. For example, we observed analysts pointing to documents and talking about their contents (e.g., 'the e. coli scare' or even using it as a proxy to stand for a person who appears in the document (e.g., 'George Prado'). Finally, documents can be involved in larger structures, such as clusters or timelines. Here they are mere parts of a greater whole. The advantage of the integration is that the low-level details including highlighted atoms continue to be available, so the contents of a document are still readable even as it participates in some larger structure. An analyst commenting on the benefit of this noted,"other tools just give me icons".

Another example of integration is the interpretation of the space. We have observed that it plays multiple roles; it is both the shoebox and the marshaling environment depending on the current intentions of the analyst. This dynamic shows the relationship between the integration of process and the integration of representation. The flexibility of integrated representations helps to facilitate an integrated process.

While some integration of representation is possible in other environments, it is clear that high-resolution space available on the large, high-resolution display provides the primary facilitation for the integration we observed.

3.3 Discussion

These studies illuminated a number of ways in which large, high-resolution spaces can support sensemaking. While the studies used nothing more than conventional desktop tools coupled with our desk-sized large, high-resolution display, we are not proposing that WordPad on a large display should replace existing sensemaking tools. In truth, WordPad and Windows XP provide minimal support for sensemaking and working spatially. The search facilities are basic, our analysts struggled to make basic annotations like highlights, and the basic window management facilities of Windows XP are not inherently spatial. There are also some basic and well known usability issues such as losing the cursor and windows and dialog boxes that appear or change in unexpected or unattended areas of the display [Robertson et al. 2005, Hutchings 2006]. However, we chose these tools because of their minimalism. We just wanted basic search and a way to display documents with a minimal amount of extraneous clutter. This allowed us to readily study core behaviors and basic representations without having to construct an entirely new set of domain specific tools for our environment.

It is also worth noting that all of our subjects were complete novices with respect to the environment. No practice or training was provided - instead, we allowed the subjects to develop strategies as the scenario required. While this did create some experimental artifacts, such as the behavior of subjects L4 and A5, the behavior that most interested us differed significantly from normal desktop use (i.e., was spatially oriented), was observed across multiple users, and showed adaptation to the environment. Despite these artifacts, these studies provided key evidence which helped to shape our theoretical framework for the use of large, high-resolution displays for sensemaking.

By creating a virtual workspace with real physical space, large, high-resolution displays offer a number of intriguing possibilities for future sensemaking tools. The main contribution of

these studies is to demonstrate that even without any special tools, the inherently spatial environment already provides support for activities typically done with physical artifacts. The study showed clear evidence of analysts using the space both as a form of rapid access external memory and as an added semantic layer in which meaning was encoded in the spatial relationships between data, documents, display, and analyst. This allowed analysts to integrate several key aspects of the sensemaking process into a flexible, visual workspace.

Chapter 4

The Impact of Physical Navigation on Sensemaking

In the previous studies, we demonstrated how the implicitly spatial environment provided by large, high-resolution displays could support the sensemaking process by leveraging the human adaptation to working with space [Kirsh 1995] and the expressive and highly flexible nature of spatial relationships [Robinson 2008b, Shipman and Marshall 1999b]. The scale of the large display allowed analysts to spatially organize ordinary text documents just as they might lay paper out on a table. While the display creates this spatial environment implicitly, there are other ways to provide an identical amount of space, including many techniques for working with large information spaces on small displays [Cockburn et al. 2009]. This raises an important question: does the way in which the space is provided to the analyst affect how it is used for sensemaking?

To address this question, we distill the issue down to one of physicality or embodiment. On the large display, the space was presented as a physical space, entirely available through the

use of physical navigation, actual physical movement on the part of the user (e.g., walking, turning, glancing) [Ball et al. 2007]. On the other end of the spectrum, we have the virtual workspace, with a conventionally sized display providing a small viewport that reveals only a portion of the available space, and access provided through virtual navigation (e.g., panning, zooming) [D. Austin Henderson and Card 1986].

In this chapter we describe a study that was designed to specifically examine how the greater embodiment provided by the large display affected the sensemaking process. In particular, our focus is on how the two navigational strategies might change the actual use of the space.

4.1 Background

The use of space for sensemaking is by no means a new idea. As we have previously noted, spatial organization plays a fundamental role in how humans manage information in the physical environment [Kirsh 1995]. In Section 2.3, we described a number of tools that were based on the spatial metaphor. Of particular relevance to this study are the various spatial hypermedia tools (e.g., [Bernstein 2003, Marshall and Frank M. Shipman 1997, Shipman et al. 2001b]), which are specifically designed for using space to marshal information. OculusÕ Sandbox is even closer to our work in that they have built on the core idea, adding specialized tools designed to help make use of the space for intelligence analysis [Wright et al. 2006]. What is notable about all of these systems is that they are explicitly designed to leverage human spatial abilities, but they are all designed to work on conventional displays, using small representations and a virtual workspace coupled with an overview+detail approach to provide the tool's space.

In Chapter 2, we described some of the prior research that has already examined some of the ways in which the physicality of the large display can affect user behavior. In particular, Ball,

Shupp, and Jakobsen all looked explicitly at display size and how it impacted physical and virtual navigation [Ball et al. 2007, Shupp et al. 2009, Rønne Jakobsen and Hornbæk 2011]. While these studies all noted the increase of physical navigation with the increase in the size of the display, the focus of the studies was on the performance of short, fairly mechanical tasks, and thus provide little insight into how this change actually impacts the user beyond potential performance differences. However, the adoption of an egocentric frame of reference when using a large display, as reported by Tan [Tan et al. 2006], and the advantages of the large display for navigating a virtual environment, as reported by Czerwinski et al. [Czerwinski et al. 2002], suggest that the difference between large, high-resolution displays and conventional displays is not purely a matter of the efficiency of physical navigation. While the advantage of having a wider field of view while navigating in a virtual environment seems to boil down to the efficiency of accessing visual landmarks physically (e.g., though a glance), rather than virtually, there seems to be a cognitive component in that the change in efficiency led to greater spatial understanding. However, while these studies suggest that the difference between physical and virtual navigation can have a cognitive effect extending beyond issues of speed and accuracy at the instance of navigation, the focus of both studies was still on the performance of short-term and fairly simple tasks.

In this study, we wish to move beyond a pure performance based approach to look at longer term behavioral changes that arise from any cognitive differences caused by the different forms of available navigation. In particular, we focus on the cognitively demanding task of sensemaking, a task for which navigation is a secondary activity serving only to support the primary task.

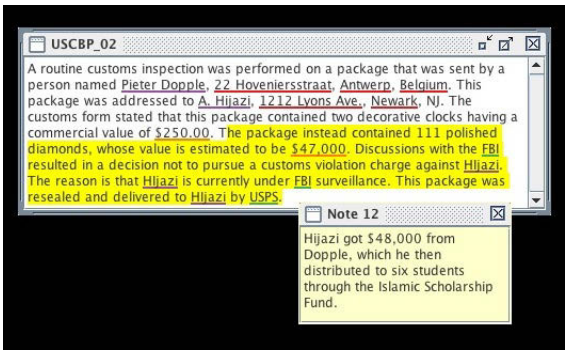

Figure 4.1: An example document showing entity underlining, a user highlight, and an attached note.

4.2 Study design

The goal of this study is to observe how the use of either physical or virtual navigation affects how a free-form spatial workspace is used during a sensemaking task. For the study, the subjects were asked to solve a synthetic intelligence analysis problem using a basic analytic tool that relies on spatial layout as the primary evidence marshalling technique. The single independent variable in this study is display size, which corresponds directly to the requirement or either physical or virtual navigation to interact with the full extent of the workspace. For the physical navigation condition, we provided the subjects with our 33 megapixel large, high-resolution display (Figure 1.2). For the virtual navigation condition, the subjects were provided with a conventional 17" LCD panel with a resolution of 1280x1024.

4.2.1 Sensemaking environment

The tool that we built for this study provides the user with a basic workspace in which full text documents can be spatially arranged to convey relationships. The familiar selection rectangle tool was provided to support the rapid movement of documents in groups, and encourage experimentation and rearrangement.

Named entity recognition was performed on the dataset using LingPipe [lin 2011], and colored underlines were used to visually highlight these in the documents (Figure 4.1). The primary purpose of this is to direct attention and to make it easier to distinguish and recognize different documents. The environment also supports two forms of annotations: highlights and notes. A highlight can be added to a document very simply by holding down a modifier key while selecting a block of text with the mouse. The notes follow the "sticky note" metaphor, and can be "stuck" to documents, following them wherever they are moved. They can also be stuck to the background as labels for regions of space, or just as freestanding text. The user can also easily change the color of notes, allowing color to be used for labeling purposes as well.

A custom file browser provides access to the documents in the dataset (Figure 4.2). The browser lists all of the available files, sorted by date, and color coded by the state of the document (unseen, open in the workspace, and seen, but no longer open). The browser also supports full text search, powered by the Lucene search engine [luc 2011].

Both conditions were provided with identically sized workspaces. On the large display, there was a direct one-to-one correspondence between pixels in the workspace and the physical pixels of the display, meaning the entire workspace was always available at the maximum level of detail. It is important to note that we purposely used fairly small (10pt) fixed size fonts so that the subjects in the large display condition could not read the text from a distance at which the display would fill their field of view. This meant that the large display subjects had to use physical navigation in order to make use of the entire workspace, while at the same time enabling the presentation of a large amount of information in the space.

The small display condition makes use of the "overview plus detail" approach to support access to the full space by providing a floating, resizable overview window (Figure 4.2).

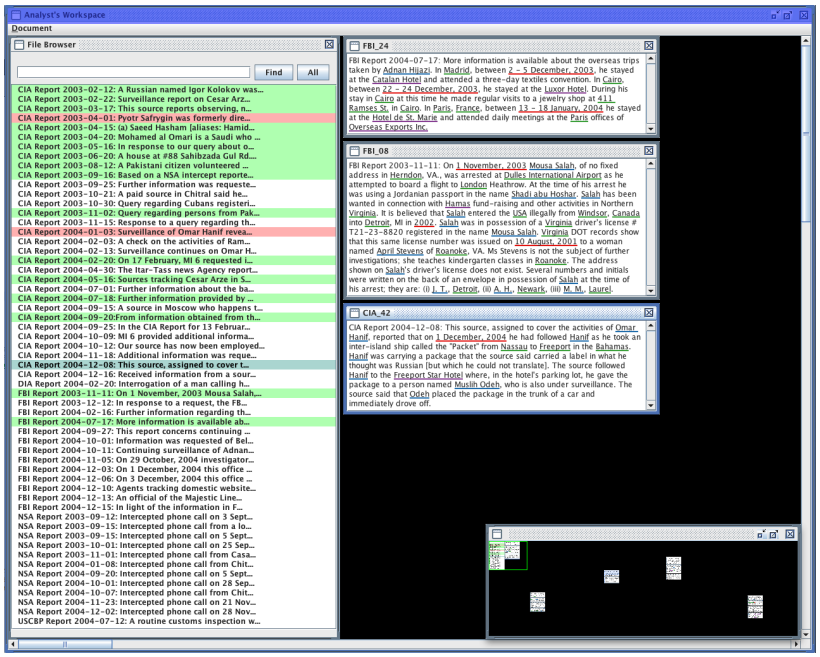

Figure 4.2: The small condition environment. A file browser is on the left, and the overview is shown in the lower right.

We chose this approach over other techniques for working with virtual workspaces because it has been shown to be the preferred technique [Cockburn et al. 2009], and it is used in many existing tools (e.g., [Bernstein 2003, Shipman et al. 2001b, Wright et al. 2006]). In our implementation, the overview is fully functional, allowing the user to select a document or documents and move them around the workspace directly from within the overview. In addition, a single click within the overview moves the viewport of the main view to the location in the workspace. This overview was available in both conditions, but open by default in the small display condition.

The viewport used in the small display also supports a number of panning techniques. Clicking and dragging the background moves the viewport around the space, similar to the move tool found in most image viewers. This implementation, however, adds a simple physics model so that if the mouse button is released in the middle of a drag, the viewport will drift to a stop, allowing the user to traverse the space faster. Dragging objects to the edge of the viewport also triggered an automatic panning, to mirror the process of dragging objects

across the workspace on the large display as closely as possible. Finally, simple scrollbars provide the conventional viewport panning interaction.

4.2.2 The task

Our subjects were asked to make sense of a collection of synthetic intelligence reports concerning the purchase and movement of bioweapons and their imminent deployment in the United States. The dataset consists of 58 documents, five of which are deliberately misleading, and approximately another ten of which only provide background information. There is no fixed starting place in this dataset, and we expected that the subjects would need to read every document in order to correctly identify the threat with enough detail and supporting evidence to prevent the attack and arrest the instigators.

Each session began with a short introduction to the analytic environment. Subjects were told that the underlying concept behind the environment was that the space could be used to express relationships between documents, but no particular strategies were discussed. The subjects were also shown samples of the kind of documents they would be working with and briefly introduced to the idea of making inferences and finding evidence to support or refute them. We then presented the subjects with the document collection and asked them to identify the underlying plot. The subjects were then allowed two hours to work with the data in the environment. After the two hours, each subject was asked to fill out a report listing the people, places, and events that they identified as significant, and to write a short description of what they thought was actually going on in the scenario. We scored the reports by assigning a point to every correctly identified person, location, or event, and subtracting a point for every incorrectly identified one (false positives). Finally, the summary was scored

out of ten points, based on how closely it adhered to the actual plot, and the amount of evidence provided.

In addition to these reports, we collected a number of other types of data. Every interaction with the tool was logged, screenshots were captured every 20 seconds, and every session was video recorded. We also provided each subject with an eye tracker, which produced a video with the target of the subject$\tilde{\text{O}}$s gaze marked in it. Each session also concluded with an informal interview about how the subject approached the problem and how he or she used the available space.

4.2.3 Demographics and recruitment

We recruited 16 subjects (six male, 10 female), eight in each condition. The subjects had a broad collection of backgrounds, being comprised of five undergraduates, nine graduate students, and two staff members, and representing a variety of disciplines including international studies, sociology, math, biomechanics, and several engineering areas.

Due to the time requirements, the study was conducted between subjects, with the condition picked for each subjects determined in alternating order based on the availability of the subject. Subjects were paid $5 an hour for their time, and a bonus of $25 for the top score within each condition was furnished to create added inducement to solve the scenario.

One subject dropped out twenty minutes before the end of the study due to a headache brought on by the eye-tracking equipment. While we do not include her score (since she did not complete a report), we do include evidence from her session in the rest of the results since she completed most of the task.

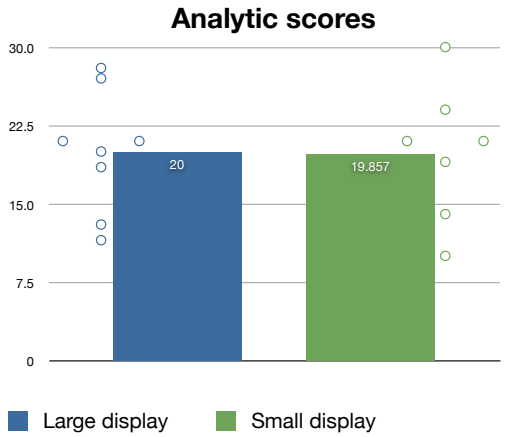

Figure 4.3: Analytic scores for the two conditions. The circles show the individual scores.

4.3 Results

The primary performance metric in this study is the score assigned to the reports. As can be seen in Figure 4.3, we found no difference between the two groups. We attribute this to the large variances found within the groups. We observed a wide spread of abilities and strategies that clearly outweighed any differences brought about by the two conditions. Given the demographics of our subjects, we also looked for differences based on gender, discipline, and age, but found no correlations. However, our primary interest in conducting this study was to observe behavioral differences, so the rest of the analysis will focus on those.

4.3.1 Use of the space

As the main purpose of the freeform spatial workspace is to encode relationships through spatial proximity, we are interested in looking at the spatial structures created by these two groups. To understand the structures created by the subjects, we used two sources of information. First, during the concluding interview, we had all of the subjects describe how they had used the workspace over the course of the investigation and then asked them to

describe how they interpreted each region of the workspace. Second, we walked through the screenshots and looked at the evolution of the structures. To try and distinguish what each structure was and the role that it played, we used a combination of the final descriptions provided by each participant, the contents of the documents and the surrounding context, as well as any notes attached to the structure by the subject.

The two primary structures that we observed being constructed were timelines and clusters, which match our earlier observations of spatial sensemaking. The organizational bases for the clusters varied, but clusters based on geographic region or person were common. Clusters also frequently had some internal organization in the form of temporal ordering, or closer proximity for closely related events. The most complex structure was a sprawling concept map developed by a large display user, in which documents were added to the structure based on some relationship to the already placed documents.

We begin our analysis by looking at the final state of the workspaces at the end of each session (Figure 4.4). That the use of space for sensemaking is a highly individualized process is clearly illustrated by these screenshots. While there are some common structural themes that run through these, they are all quite different. To further illustrate this, we will describe some of the more interesting examples.

In Figure 4.4(b), the subject has created a loop documents. This subject made no attempt to use the space available, and instead just opened documents and let them lie where they were, moving the viewport whenever the view got too crowded. The loop was formed as the viewport arced down to the bottom of the display and then wrapped around to the left. In essence, this subject quite literally worked himself into a corner. In contrast, Figure 4.4(l) shows one of the only small display subjects to attempt to make use of the whole available workspace. He used a structure based roughly on geographic relationships, with

important documents corralled to the right. Most of the others made small clusters of related documents, but there were no other approaches that assigned any overall structure to the space.

On the large display, our highest scoring subject produced the aforementioned concept map that can be seen in Figure 4.4(e). Two of the subjects created timelines of important events and surrounded them with clusters based on important people, organizations, and locations (Figure 4.4(g & o)). Figure 4.4(k), on the other hand, split the space, with important people clustered above, and arranged by relationships, and events and places clustered below. The workspace shown in Figure 4.4(c) is a curious one in that the subject made use of an overall geographic layout for the space, but notes have replaced all the documents. As she read each document, she replaced it with a note containing a couple of keywords she thought might be important and then closed the document. Unfortunately, she did not know what would be important and at the end only had a collection of names and places without any of the connecting events that tied them all together. Her low score provides compelling evidence for the benefit of easy access to detailed information (i.e., the actual text of the documents), and the need for evidence marshalling techniques to evolve with understanding of the problem space.

Moving beyond the individual characteristics of these workspaces, we can get a sense that the subjects using the large display were considering the entire space, and developing their organizational strategies accordingly. While these silhouettes hide the underlying details of the documents, there is an appearance of coherency across the entire workspace. In contrast, most of the workspaces of the small display subjects seem cramped and lacking of an overall structure or coherence. In the next few sections we will break down some of the components of these layouts to provide deeper insights into how these two groups differed.

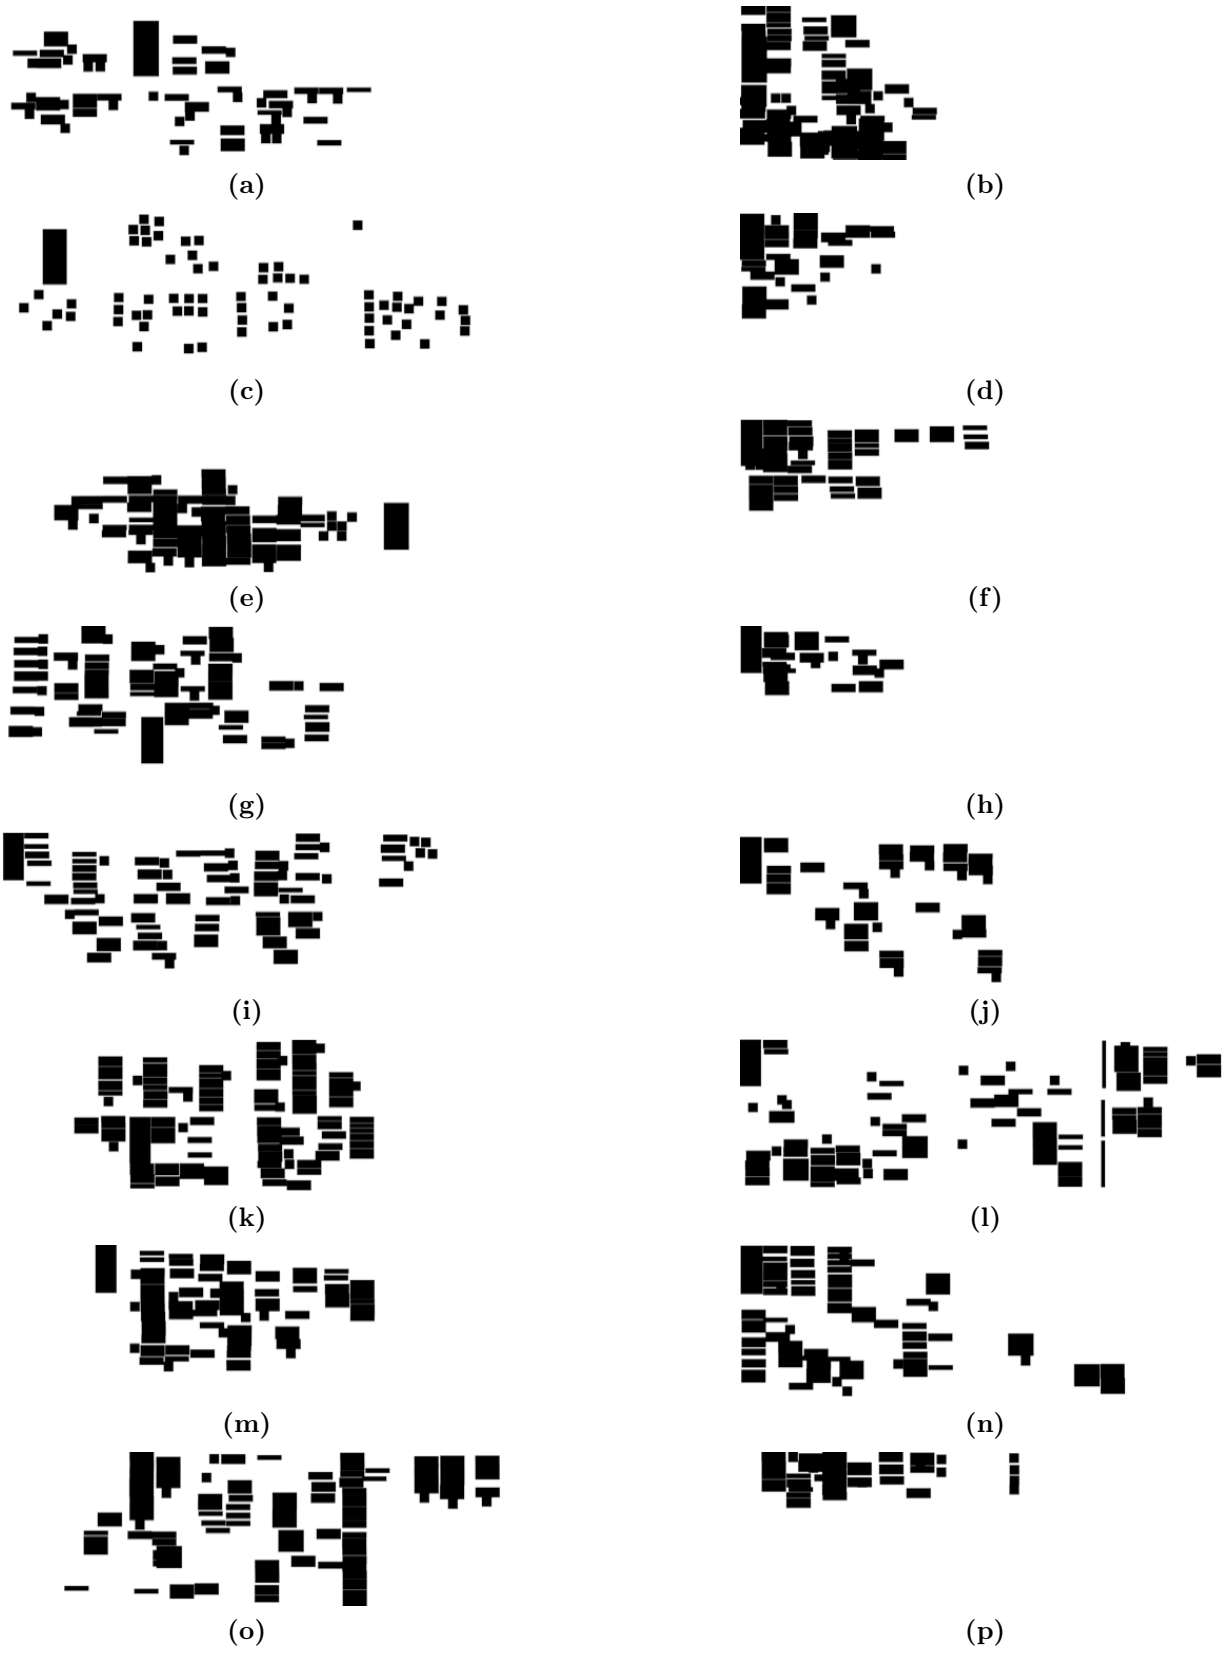

Figure 4.4: Simplified silhouette views of the final workspace state for all users. The large display condition is represented in figures a,c,e,g,i,k,m,o and the small display condition is represented by figures b,d,f,h,j,l,n,p

Table 4.1: Summary of the results of the quantitative measures

Measure	Large display	Small display	p value
display coverage	18.53%	13.12%	0.061
display span	52.26%	34.39%	0.095
# of structures	15.75	9	0.004
# of documents in the largest structure	22.375	11.75	0.095
# of complex structures	4.35	0.875	0.065
% of structures that are complex	29.7%	7.7%	0.086
# of structural events	96.875	45.25	0.011
# of structural events per structure	6.63	4.87	0.271

4.3.2 Quantitative metrics

For all of our measures, we use an unpaired, two-tailed Student's t-test to determine significance. Given the exploratory nature of this study, we set the threshold of significance at $p < 0.1$. As such, we do not wish to overemphasize any particular result, but rather wish to demonstrate an overall pattern of behavioral change. Table 4.1 provides a summary of these quantitative results.

Workspace utilization

The most obvious difference is that the large display condition used more of the workspace. To determine this we looked at two different measures of space utilization: coverage and span. By coverage, we mean the amount of the workspace actually covered by some visual artifact (document, note, browser, or overview). Span, on the other hand, is an attempt to quantify how spread out across the workspace the visual artifacts are. We calculate this by computing the convex hull around the artifacts and then using the percentage of the workspace that this consumes as the span. For both measures, we show a significant increase for the large display condition. The large display group covered 5.41% more of the display

and spanned across 17.87% more of the display. While the former is a fairly small difference, it is interesting that there is any difference at all, since both groups were dealing with the same number of documents. This difference implies that the smaller display group closed more documents and had overlapping documents more often (though neither of these show a significant difference when examined individually).

Spatial structures

In looking at the structures created by the subjects, the first thing we observe is that the large display group created 75% more structures over the course of the investigation. This group also created larger structures, with the average size of the largest structure being 90% larger for the large display group.

We also looked at the complexity of the structures, since complex structures should be indicative of higher-order schemas being externalized into the space. We tallied up the number of structures that showed some form of internal structure, be it temporal ordering, or some form of hierarchical construction. The large display group produced 385.7% more complex structures than the small display group. If we take into account the larger number of structures created by the large display group, we can break this down to find that 29.7% of the structures created by the large display group were complex, as opposed to 7.7% of the structures created by the small display group.

We also found that the large group produced 114% more structure changing events. We consider a structure changing events to be the movement of a document or documents that changed the size of a structure or the location in the space. Given the difference in the number of structures, this is perhaps not surprising, and there is no significant difference when we look at the number of events per structure.

So, despite a wide variety of strategies being employed by our subjects, we have shown that the different displays have an impact on the structural complexity of the workspace. On the large display, we find that the workspace is overall more dynamic, with not just more structures being created, but also larger and more complex ones as well.

4.3.3 Investigative approaches

We observed a wide variety of strategies being applied to this problem. However, for the most part, all of the subjects generally followed a strategy Kang et al referred to as "overview, filter, and detail" [Kang et al. 2009], though none of our subjects adhered to it slavishly. In essence, this involves scanning through all of the documents fairly quickly, and then revisiting interesting documents and keywords with greater scrutiny.

Since the subjects were primed by the initial training, almost all of them, in both groups, began by using the space to store documents of interest. In most instances this meant clustering documents based on some surface feature that seemed distinctive, usually location, but sometimes person or activity. An interesting difference between the two groups that manifested itself fairly early was that the large display group typically categorized all documents as soon as they were read. The small display group, however, frequently left documents where they opened unless the document seemed particularly interesting, was connected to other documents that were already categorized, or was related to another recently read document that had not yet been categorized. The result of this was that the small display users frequently formed piles of overlapping documents in the viewport.

Following the initial scan through the documents, there were two primary techniques that were followed. The first was to visit each cluster or structure in the space, rereading the

documents found there, and then rearranging them into new structures based on the fuller knowledge of the dataset. In other words, these subjects were applying incremental formalism. We observed this behavior to some extent in most of the large display subjects, while only one of the small display subjects really committed to this approach.

The other technique was to employ the search tool to follow up promising keywords. These two techniques are not mutually exclusive, but many of the small display users began to use this technique exclusively, searching and opening new copies of the document, rarely revisiting any structures that they made earlier. We can see evidence of this by looking at the how the small display users moved their viewport around the workspace. Figure 4.5 shows a collection of heatmaps that were generated by tracking where in the workspace the viewport was positioned. We can see by the white hotspots that in most instances there was a single point of focus where the subject spent the majority of his or her time.

The large display subjects performed just as many searches on average, but they used the search results differently. They would open new copies of the documents in the search results, but only to see what they looked like. They would then scan quickly over the workspace to find the original copy, closing the new copy when they had found it. This behavior clearly suggests that some form of visual search tool that highlights location is needed for similar environments. However, we would also submit that this indicates a different regard for the documents open in the workspace. The subjects clearly felt that the spatial location or the documents in the surrounding region added contextual information not available in the raw words that were already fully available in the second copy that they opened from the search results.

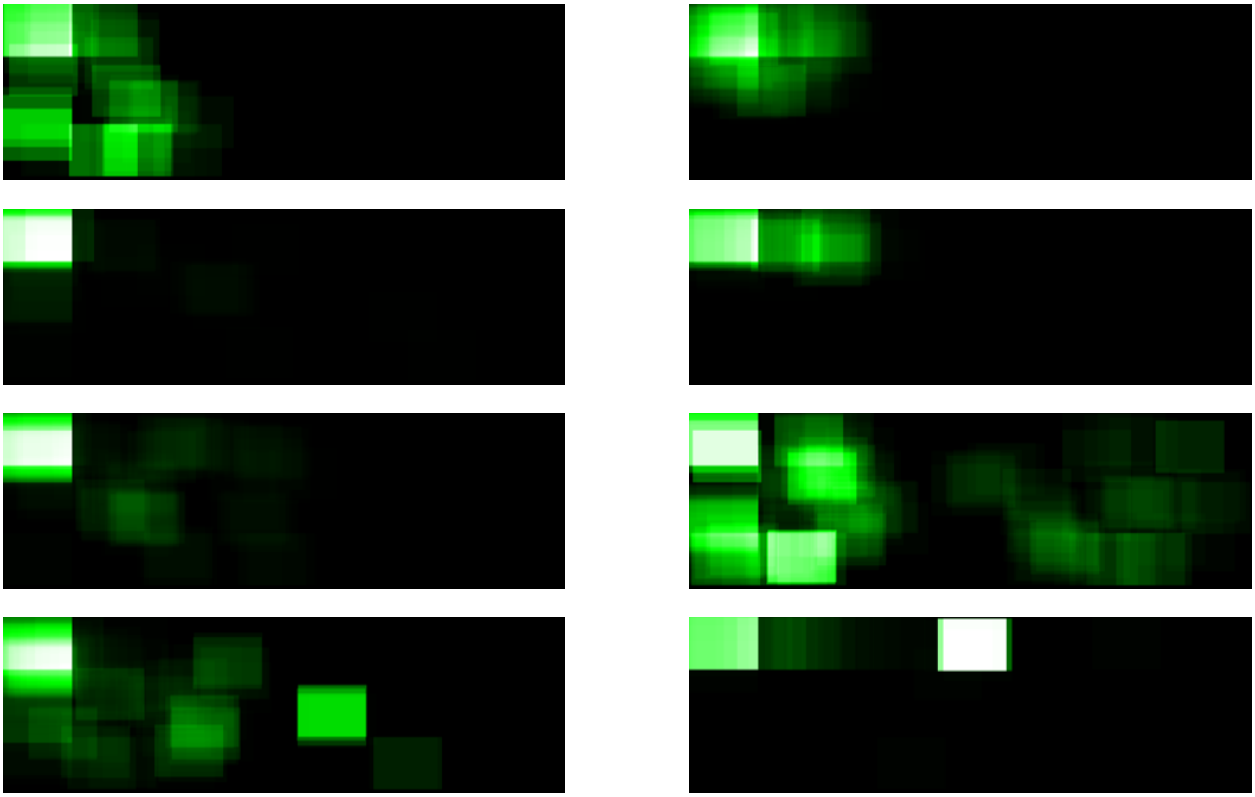

Figure 4.5: Heatmaps showing the locations of the viewport of the course of the sessions for the eight small display subjects.

4.3.4 Revisiting documents

Looking at the eye tracking data, we find evidence that the large display group revisited documents more often. In Figure 4.6, we provide a visualization of representative users in both groups. Each horizontal row represents a document or other object (note or tool) in the collection, sorted by type and then by title. As can be seen, each subject opened documents in approximately the same order, causing the initial diagonal. However, to the right of the diagonal, we can see that the large display user refers back to open documents far more often than the small display user. The thin, almost vertical lines that appear in the large display userÕs visualization illustrate moments when the subject was scanning the workspace, either looking for a particular document or just refreshing her sense of the workspace.

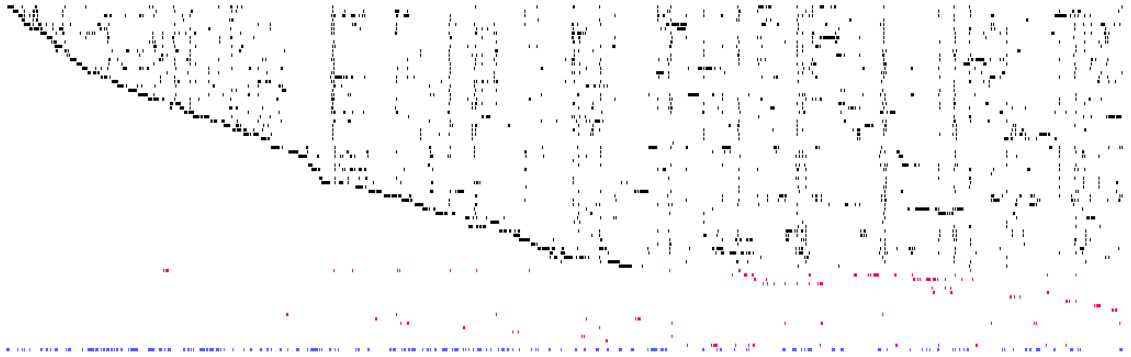

(a) large display subject

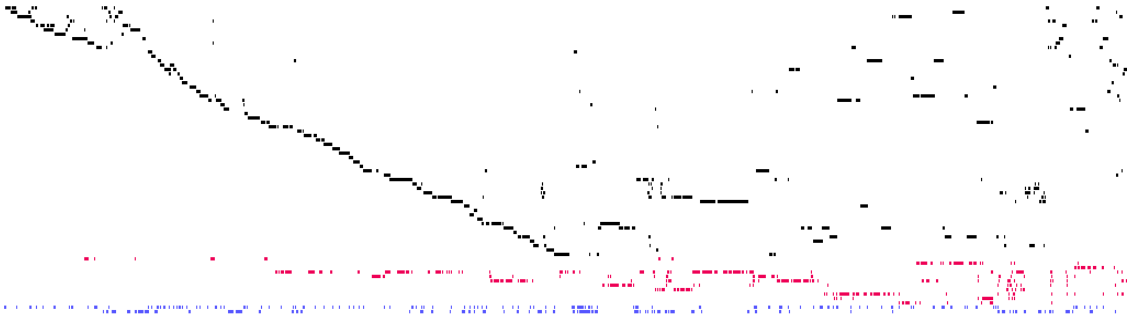

(b) small display subject

Figure 4.6: Visualization of eyetracking data for two subjects. The x-axis is time and each row represents the document (black), note (red), or tool (blue) at which the subject is looking.

Unfortunately, due to some technical issues with the eye trackers and the difficulty in coding the results, we only have quantitative tracking data for a few of our subjects, so we cannot demonstrate a statistical difference between the groups. However, this behavior is apparent in all of the sessions we were able to code. This is also supported by what we observed in the movement of the viewports on the small display and the search behavior.

It is fairly clear that the increased access to the documents is attributable to the efficiency difference between physical and virtual navigation. While there is insufficient evidence to tie the increased access to any observed behavioral changes, it certainly reduced the amount of information that the subjects had to hold in their heads, while simultaneously providing the opportunity to internalize it through constant exposure. The more complex structures that we observed might also be attributable the spatial context provided by greater awareness of

the surrounding documents. It is clear that this is a phenomenon that requires more study to fully appreciate the cognitive implications.

4.4 Discussion

The assorted evidence that we collected indicates that there was a significant behavioral difference between the groups with respect to the way they viewed and used the available space for sensemaking.

4.4.1 Externalization

We claim that the large display subjects viewed the space as a more coherent whole and that this lead to more effective externalization into the space. Note that in this instance we are referring to effectiveness of the environment in supporting externalizing, rather than the effectiveness of the actual externalizations for sensemaking.

As we can see from the reported results, the large display subjects externalized more information into the environment, in the form of additional structures as well as more complex structures, which encoded more information than simple clusters. The treatment of the environment workspace as a coherent space allowed this group to treat the dataset as a single connected whole, providing the opportunities to link together clusters of related information. As we observed when discussing search strategies, structures and placement seemed more important to the large display group than to the small display group, which rarely seemed to feel the need to relocate a document in its spatial context. We also note that in discussions at the conclusion of the sessions, one of the questions we asked was "how would you

approach this problem if you had to do it again?" The majority of the large display subjects replied with some variation on suggestions of alternative organizational strategies based on what they learned. Among the small display subjects who had pursued some form of spatial organization initially, however, the general response was that they wish they had abandoned spatial organization earlier and concentrated on taking notes instead. This suggests that the large display provided far stronger support for spatial organization. It is, of course, worth asking if this difference is really important given the analytic results. We would argue that it is. The wide variance in the scores can be attributed to many factors, but the primary one is inexperience, both with the domain and the use of space for sensemaking. The difference between a good solution and a poor one could sometimes hang on which document first seemed suspicious to the subject.

It is worth acknowledging that the alternate organizational strategies we mentioned above may be an indication of the perceived importance of spatial organization, but it also implies that the organizations adopted by the subjects were not optimal. While one of the benefits of space as a structural medium is the flexibility to restructure with the introduction of new material, there are some primary strategies suggested by more domain experience that would have helped these subjects.

The key point here is that this study has demonstrated that the physicality afforded by the large, high-resolution display led to more information about the relationships between people, locations, and events to be externalized spatially than occurred with the virtual workspaces.

4.4.2 Physical navigation

Taking a step back, the most important point that this study makes is that there are cognitive differences between physical and virtual navigation. Earlier research, as we discussed

previously, has shown the use of physical navigation has positive performance benefits for visualization tasks [Ball et al. 2007, Shupp et al. 2009]. However, given the mechanical nature of the tasks, an argument could be made that the benefits of physical navigation were attributable to the mechanical efficiency or directness of pure physical movement that has been unmediated by an interaction device. That same argument cannot be made in this case, however. Our results do not speak to raw speed, but instead demonstrate clear behavioral differences over a prolonged period.

There are a number of potential explanations for why physical navigation would influence perception and behavior. We postulate that these are rooted in the greater support for embodied resources afforded by the display environment.

For example, we can return to O'Regan's postulate that the visual system works by using the physical world as a form of on-demand external memory, and that the rich, cohesive visual environment that we perceive is actually a form of illusion created by the immediate availability of external information [O'Regan 1992]. It is possible that the effect we are observing is an extension of this property of the visual system. In other words, the large display subjects are making use of core properties of the visual system to build up a coherent sense of the space based on constant and immediate access. The small display subjects, on the other hand, do not have the directness of visual access. Instead, they need to maintain a mental mapping between the low detail overview, an internal representation that interprets it, and the detailed view. Another possibility is offered up by Patten and Ishii's suggestion that being able to relate to objects in an egocentric frame of reference adds additional semantics to the spatial location [Patten and Ishii 2000], and Tan's demonstration that physical large displays bias users into adopting an egocentric frame of reference [Tan 2004]. It is possible that the physical nature of the large display allowed the subjects to relate to the spatial

locations better, making it easier to distinguish different location in the space and thus easier to use to express relationships through spatial relationships.

In truth, the differences are probably created by a combination of these factors and many others as well. However, it seems clear that the embodied nature of physical navigation taps into cognitive resources that are unavailable through the use of virtual navigation. This in turn makes it easier for the user to make use of the available space for externalization. Locations in the space are more easily discernible and can be related to the context of the rest of the userÕs environment. The coherence of the space as maintained by quick eye movements and optical flow cues means that the user can manage larger, and more complex structures.

4.5 Potential issues

There are, of course, some issues and open questions raised by this study. One potential confound is the presence of bezels on the large display. The bezels divide the large display up into a series of "bins" that arguably provide additional structure to the large display workspace, which was not provided in the virtual workspace. The effect was certainly present, but we would argue that looking at the final workspace states, there is fairly little evidence that the bezels played a strong role in the organization of the space. We examined these screenshots with an overlay showing the location of the bezels, and found minimal use of the bezels as an organizational tool. The only real examples can be seen in Figures 4.4(a) and 4.4(k), where the horizontal line separating the two rows of displays clearly separated the structures into a top and bottom. However, even in these instances, the two regions did not seem to assume any additional role or meaning, it was just a dividing line that structures did not cross.

It could also be argued that the environment we developed favored the large display, as the full text documents can fill the small display rapidly, without much opportunity for the interesting construction of spatial structures within the confines of the viewport. However, part of the point was to juxtapose virtual and physical navigation, and any interaction that was required to access the contents of a document would have added virtual navigation to both conditions. More importantly, the very fact that the environment could favor the physical space condition only serves to strengthen the assertion that a large, high-resolution display cannot be simply replaced with a virtual workspace on a conventional display.

The inexperience of our subjects is another key issue. This is a challenge when doing any research in this domain, or indeed any domain that typically requires a level of expertise. Given the time requirements to do any kind of real sensemaking it is difficult to either find trained personal or to train novices to an appropriate level. Ideally, future exploration into this issue will involve more skilled subjects with the goal of finding real performance difference between the two environments.

A more serious concern is the cognitive effect of managing large quantities of information spatially. This is, of course, more of a general issue with the approach than a problem with the study. It seems clear that cognitive cycles must be used to maintain spatial understanding and to perform spatial organization. In about half of our large display subjects we observed some evidence of what we might call "organizational thrashing", where documents were being organized and reorganized but there was not any real understanding being developed by the process. It is possible that the large display subjects actually hurt their scores by spending too much time reorganizing rather than actually reading and making new connections.

We continue to pursue this particular approach to sensemaking because we feel that the cognitive overhead of managing the space is easily outweighed by the cognitive overhead of

managing large information collections internally, or with more linear or structured externalization techniques. However, this is clearly an area that needs more research.

The final issue that we are frequently asked about is the issue of scalability. Virtual navigation allows the user to work with workspaces of any size, while strict adherence to physical navigation imposes fairly clear boundaries on the amount of information that can be displayed. Our initial response is that we have demonstrated a clear difference in how space on the large display is used, and that difference will not disappear if the workspace now grows beyond the bounds of the display, or some other data aggregation technique is employed, requiring a mix of virtual and physical navigation. Our interest is in demonstrating the fundamental difference between the environments rather than proposing a complete sensemaking environment. However, as indicated by Jakobsen and Hornbæk [Rønne Jakobsen and Hornbæk 2011], it seems clear that future research will need to look into how the integration of aggregation techniques and virtual navigation affects the large, high-resolution display environment.

4.6 Conclusion

Physical navigation is fundamentally more embodied than virtual navigation, which requires technological mediation and internal mappings to maintain spatial understanding. As such, large, high-resolution displays provide more opportunities to leverage embodied resources than conventional displays, which must rely on virtual navigation. In this study, we have shown that the opportunity to use physical navigation led to the development of more effective externalization. Users of the large display made use of more of the available workspace, treated the workspace as a more coherent whole, and created larger, more complex structures.

This change did not result in significantly better analytic performance, but we attribute this primarily to the wide variance between the subjects. The important point, however, is that there was a behavioral shift that indicates that the physical workspace provides a different set of cognitive affordances.

The virtual workspace is still a useful technique for interacting with large information spaces, but this study has demonstrated that all space is not equal. Physical navigation cannot simply be replaced with virtual navigation techniques with the expectation that the userÕs spatial abilities and thus the facility to leverage the space to externalize information will carry through identically. Large, high-resolution displays provide a fundamental different environment than virtual workspaces. The strong support they provide for large quantities of detailed information and human spatial abilities make them a compelling platform for future sensemaking tool development.

Chapter 5

Analyst's Workspace

Analyst's Workspace (AW) is a new sensemaking environment, built with the express goal of providing an environment that unifies the activities of foraging and synthesis into a single investigative thread. The problem that AW is trying to address is a disconnect that we have observed in many tools with respect to synthesis. Either synthesis is left entirely in the analyst's head, which can lead to potential problems with cognitive bias and an inability to communicate how a particular conclusion was reached [Heuer 1999], or it is done separately from the foraging activities, eschewing the tight connection between the two. This is not overly surprising, given that foraging is computationally more tractable. Searching, filtering, visualizing and other information retrieval techniques for dealing with large collections of information are all clear ways in which we can apply technology to the foraging process. The synthesis process is much more difficult to support directly. It requires expert knowledge and even if we could computationally schematize, there is a question of the validity of a schema that has not been internalized (i.e., understanding can only be promoted, not granted). AW seeks to address this problem by providing lightweight tools that encourage the analyst to

externalize aspects of the synthesis process into the environment without interrupting the investigative flow.

The design of AW is based primarily on the work described in the earlier chapters, and was developed explicitly for use on a large, high-resolution display. As we observed in the analytic studies of Chapter 3, the scale of the display encouraged the adoption of a spatial approach to managing information. In the studies, the subjects worked with full text documents, using them simultaneously as primary source material, context for extracted information, and representations of broader themes – thus supporting both foraging and synthesis with the same representational artifacts.

While the environment provided basic support for analysis, our analysts encountered numerous issues. Modern windowed systems support the movement of individual windows, but they are not optimized for large-scale organization, and our analysts reported making compromises in their spatial organizations due to the difficulty of making changes that would affect more than one or two documents. In addition, there is no easy way to annotate documents or make connections between them, and the basic search provided by the operating system for searching across the set of documents is a somewhat limited foraging tool.

Analyst's Workspace was designed to address these issues, leveraging the observed advantages of the spatial workspace afforded by the large display into a robust tool supporting an integrated approach to foraging and synthesis. A secondary aim of this work is to explore some of the design space surrounding the development of analytic tools for large, high-resolution displays. In this chapter, after a discussion of related work, we describe the basic operation of AW, followed by the rationale and motivations behind its design.

5.1 Related Work

In Section 2.3 we surveyed a collection of tools designed to aid sensemaking. Here, we will reexamine a few of them with respect to the range of activities that they support across the sensemaking process.

In the early stages of foraging process, a key activity is gaining an overview of a collection of data and filtering it down to a more manageable size. IN-SPIRE is an example of a tool designed specifically to address this activity. While it provides some sophisticated query and time-series capabilities, it is predominantly a tool for text clustering, using visualization to highlight important topics or themes in the data set [Wise et al. 1995]. As such, it is fundamentally a tool concerned with activities early in the foraging process, where the concern is primarily focused on the development of a "shoebox" of information to pursue further. This is a phase of the sensemaking process that is not directly supported by any of the other analytic tools we will examine, including AW. Other tools seem to be directed at more focused tasks or datasets that do not require the analyst to start from a completely anonymous collection of documents without a clear goal or target.

The link analysis tools, such as Analyst's Notebook [Ana 2011], Sentinel Visualizer [sen 2011], and Palantir Government [Pal 2011] focus predominantly on identifying and visualizing connections between entities. The primary activity supported by these tools is exploration – chaining from entity to entity to find interesting patterns of connections (essentially the "Search" arrows in Figure 2). However, as we noted earlier, the context of the data can be obscured due to the lack of emphasis on primary documents, leading them to be used more to manually organize the results of an investigation as part of writing a report.

Jigsaw is another tool that is fundamentally concerned with entities and the connections between them [Stasko et al. 2007]. The underlying motivation for Jigsaw, like AW, is to

provide support for the analyst – not to replace the analyst with algorithmic processes. Unlike the link analysis tools, the focus is on directing the analyst to the important reports that still need to be actually read. Jigsaw provides the analyst with a wide variety of visualizations that can be used to illuminate different aspects of the data set and focuses the analyst's exploration. Jigsaw has recently incorporated a "Tablet" tool, which provides a freeform space for schematizing and otherwise recording important information. However, Jigsaw is fundamentally an investigation tool, and best supports foraging activities.

Entity Workspace directly addresses the very tail end of the foraging loop [Bier et al. 2006]. The core purpose of Entity Workspace is to serve as an evidence marshalling tool, allowing the analyst to rapidly make and record connections between entities as they are encountered in the text. The product produced by Entity Workspace is an evidence file, which can then serve as the basis of schema building and synthesis. However, like AW, Entity Workspace is fundamentally concerned with an integrated investigation and allowing lightweight, unspecified connections to facilitate analytic externalization. In fact, we feel that the entity composition tools are very complimentary to the core goals of AW and would make a compelling addition to the AW toolset.

Occulus' GeoTime is a similar tool with a different focus. Whereas Entity Workspace helps the analyst to marshal information about entities, GeoTime helps to marshal geographic and temporal information into a single visualization [Kapler and Wright 2005]. As with the link analysis tools visualizing these patterns can help drive the development of schemas, but it still leaves the process completely in the hands of the analyst. That said, the later introduction of stories added a form of externalization to help the analyst manage schemas and hypotheses, bringing it closer to the goals of AW [Eccles et al. 2007].

As noted earlier, Sandbox's spatial workspace is conceptually quite similar to AW, though it provides a larger collection of analytic structuring tools and can export into a number of

report formats [Wright et al. 2006]. Sandbox also part of the nSpace suite, and is integrated with the Trist system, which provides a powerful collection of search tools [Jonker et al. 2005], permitting the pair to cover a wide range of sensemaking activities from basic search through reporting. However, because the foraging is relegated to a separate application, Sandbox cannot offer quite the same tight level of integration offered by AW. In addition, as illustrated in the previous chapter, the use of real physical space allows AW to include more information (such as full text documents), and more distinct locations, providing greater context for semantic structures and thus more expressive externalization of information.

5.2 System Overview

At its simplest, Analyst's Workspace is a spatial environment for sensemaking. The analyst can use AW to explore a collection of documents, opening them in the space and then arranging them into meaningful patterns as part of the sensemaking process. In this section, we will provide a high-level overview of the structure and capabilities of the tool.

5.2.1 Physical Environment

Unlike most analytic tools, the physical environment in which AW is used can be considered a core component of the overall system. AW was specifically designed to run on the large, high-resolution display seen in Figure 1.2. While this display is certainly large enough to be used in a collaborative setting, it is important to note that this setup is configured specifically for individual use, and correspondingly, AW was developed as a single user application.

While AW could be modified to run on a smaller display using a virtual workspace to replace the real physical space, the results of the study described in Chapter 4 demonstrate

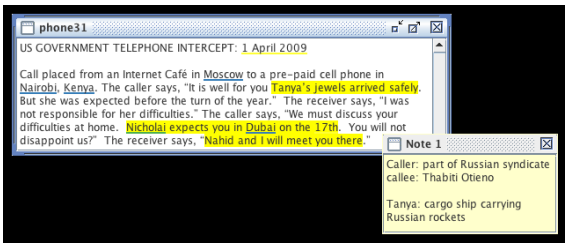

Figure 5.1: AW document showing entity underlining, user highlights and an attached "sticky note"

that the utility of the space for the externalization of the synthesis process through spatial organization would be compromised. As such, the display is an inseparable piece of AW's design, and the design decisions made in the course of developing AW were made purely to support the fully physical space afforded by the display.

5.2.2 Documents

The primary technique for schematizing information in Analyst's Workspace is the spatial organization of documents. Unlike many tools, documents are not represented in abstract or iconified form, but as actual frames of text (Figure 5.1). This allows the analyst to continue to access the details of the information, even as the documents are arranged to represent higher-level connections and narrative relationships. The documents are resizable, and a conventional scroll pane permits some text to be hidden, but by default, documents are automatically sized to display the full text.

The general focus of AW is on working with reasonably short, free-form text documents containing a few key facts and a couple of paragraphs. This structure matches the format of most news articles and intelligence field reports, as derived from the VAST challenge datasets [Costello et al. 2009] and various intelligence training exercises that have been made available to us. Longer documents can certainly be analyzed in Analyst's Workspace, but documents that contain only a couple of key points are more easily abstracted to a single concept,

making them conceptually easier to use as building blocks in the construction of narratives or intermediate schemas, a lesson we learned from the more limited use of space we observed in Chapter 3 when the consolidated document was employed.

Since documents play both the role of source material and semantic building block, AW allows documents to be cloned in the space so that they can be used in multiple locations. The cloned documents share annotations and are connected by visual links. Obviously, while this facility can be used for all documents, it is particularly useful for documents containing multiple events or concepts.

In addition to text, Analyst's Workspace also supports image documents. To support search, we use specially formatted index files containing a link to the associated image and any relevant metadata that might help the analyst connect the image to other parts of the investigation.

Data ingest is handled in a couple of ways. AW can load single documents or whole directories of raw text files. For additional functionality, however, AW will also accept specially formatted XML documents that bundle metadata such as report dates, titles, and associated entities along with the raw text.

A file browser can be called up on demand within the tool to access all of the currently loaded documents (Figure 5.2). This provides a straightforward listing of all known documents, allowing the analyst to open any document of interest. The browser also provides simple controls to sort the files by title, the name of the original files, or by date. Color-coding indicating document state is used to provide an overview of the state of the investigation. The blue background indicates that the document is currently open in the workspace, while gray text means that the document has been visited, but is now closed. If the document is uncolored, it means that the analyst has not yet viewed the contents. This coding allows the

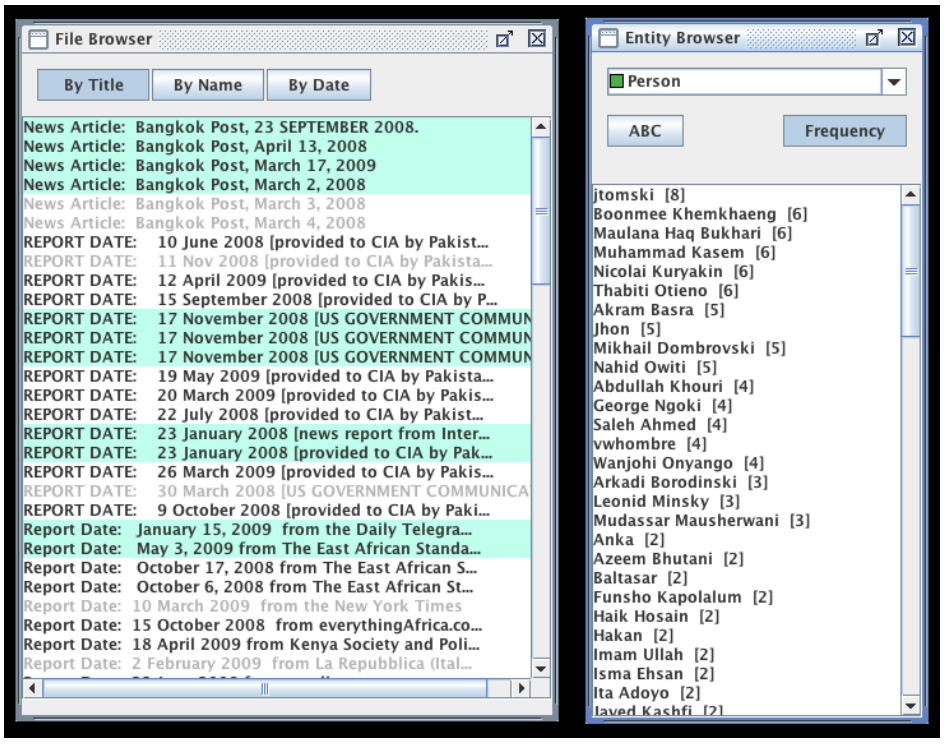

Figure 5.2: File Browser [left] and Entity Browser [right]

analyst to rapidly assess how thoroughly the investigation has covered the current dataset. This color scheme is also used consistently throughout the interface for all document lists.

5.2.3 Entities

Entities, such as people, places, organizations, dates, etc., play an important role in the kind of investigations carried out by intelligence analysts. Like most of the previously mentioned analytic tools, Analyst's Workspace relies on entities as core elements of its analytic approach.

Initial entity identification is performed using LingPipe [lin 2011]. However, since automatic identification is far from perfect, manual entity identification and editing is fully supported. The analyst can easily highlight text and use contextual menus to assign it to an existing

entity category or to a new entity type. Entities can also be quickly aliased to other entities in a similar fashion.

Entities are first-class visual objects in AW. They are identified in documents, but they can be opened directly into the workspace where they appear as a labeled list of all of the documents in which it appears (Figure 5.3). Entities can also be iconified. This alternative representation is intended to make it easier for the analyst to cognitively relate to the entity object as the entity, rather than a list of documents, thus making it easier to use the entity as a label, or to create spatial structures representing entity relationships. The iconified form is also intended to make it easier for the analyst to rapidly find and identify all of the important entities of a certain type (e.g., current suspects and their connections).

To aid the analyst in finding connections within the data set, AW maintains a co-occurrence graph of all of the entities. We take the same simple approach to this as Jigsaw, and consider two entities to co-occur if they both appear in the same document [Stasko et al. 2007]. The analyst gets access to these connections visually through the use of colored lines that are drawn between entities displayed in the workspace. Selecting one of these links reveals the number of documents shared by the two entities and double-clicking the link opens a list of the shared documents.

Analyst's Workspace provides the analyst with access to all of the entities identified by the system through on-demand entity browsers (Figure 5.2). The browsers allow the analyst to see all of the entities separated out by category. These lists can be sorted by either alphabetical order or by frequency (i.e., the number of documents in which the entity appears).

5.2.4 In-place annotation

Analyst's Workspace provides the analyst with two forms of in-place annotation: notes and text highlighting (Figure 5.1). The notes resemble conventional "sticky notes", and support free-form text. The color is customizable, allowing the analyst to use them as labels, or to maintain a variety of different kinds of notes in the workspace, assigning different semantics to different colors. Like the documents, notes can be placed anywhere in the workspace and moved freely. However, notes also conform to the "sticky note" metaphor and can be "stuck" to another document by dragging the top edge of the note onto the document. The note will then follow the document whenever the document is moved.

As we observed in the earlier study, textual highlights are an important analytic tool when working with textual documents. They allow the analyst to extract information from a document in a way that keeps it within the original context. It also serves to make documents more visually distinct and easier to recognize. For these reasons, and the need to highlight text frequently while reading, AW makes highlighting very easy. The analyst only has to hold down a modifier key on the keyboard while selecting text and it will be highlighted.

Highlights are atomic elements that are bundled directly with the documents that they modify. If additional copies of a document are opened in the space, they will all share the same highlighting. Highlights can also be selected, and then deleted or dragged out of the document. If the highlight is dragged out into the workspace, the system produces a new note containing the text. If it is dragged on to an existing note, the text is added at the drop location.

5.2.5 Following a trail

Analyst's Workspace supports two primary techniques for pursuing an investigation. The primary technique is to follow an entity-directed investigation technique. All of the extracted entities are identified in the documents through the use of colored underlines. Graphically, this helps analysts to skim documents and rapidly get a sense of the key players and locations. However, the main use of the underlined entities is to allow in-place pursuit of interesting leads.

When the analyst hovers the cursor over an entity, information scent in the form of a tooltip appears that lists any aliases of the entity as well as the number of documents in which the entity appears. If an entity seems interesting, clicking on it opens the entity object into the workspace.

In Figure 5.3, we see an example document and an open entity. The analyst has identified Thabiti Otieno as a person of interest, and the tooltip confirms that he appears in several documents. The entity is open below the document, showing a list of the other documents in which Otieno appears. We can also see the visual links connecting the entity to the original document as well as several others. In this instance, we see that there are two other open documents which reference Otieno, indicating that there is a previously unexplored connection from the current document back to documents found earlier in the investigation, increasing the potential significance of this entity.

While entity-based exploration is anticipated to be the primary investigative technique, AW also supports full-text search across the data set. The actual search functionality is implemented using Apache Lucene [luc 2011], which provides a full set of search operators to the user.

To invoke the search tool, the analyst can use the global context menu or a keyboard shortcut. To speed entry, the tool will be automatically populated with the last textual selection made by the analyst, so search can be performed by selecting an interesting phrase, invoking the search tool, and then tapping the ENTER key to perform the search.

Search results are persistent objects that bear a close resemblance to entity objects. There are two primary differences however. First, provenance for the search is provided in the form of a directed link from the originating document. Second, AW highlights the search term in all open documents that contain it, helping the analyst to rapidly assess each document's value with respect to the analyst's current interest.

One of the goals of these techniques is that the analyst should be able to find more information about an entity or phrase immediately without leaving the context of the document. In the course of reading a document, an analyst may even collect a number of leads, which will all then be open in the space providing a visual reminder of which leads to pursue as well as the means of pursuing them. We feel that this approach will also influence the analyst into adopting an approach to the data that Kang et al. labeled as "Find a clue, follow the trail (FCFT)" [Kang et al. 2009]. This technique is not oriented on narrowing in on a solution from an overview of the entire collection, but instead cuts through the dataset based on leads in the data, making it an efficient approach when there are many irrelevant documents in the dataset as these will not be uncovered by the process.

5.2.6 Visual links

Visual links play another important role in the interface as they allow the analyst to see connections that are not represented by spatial proximity. These links are drawn automatically

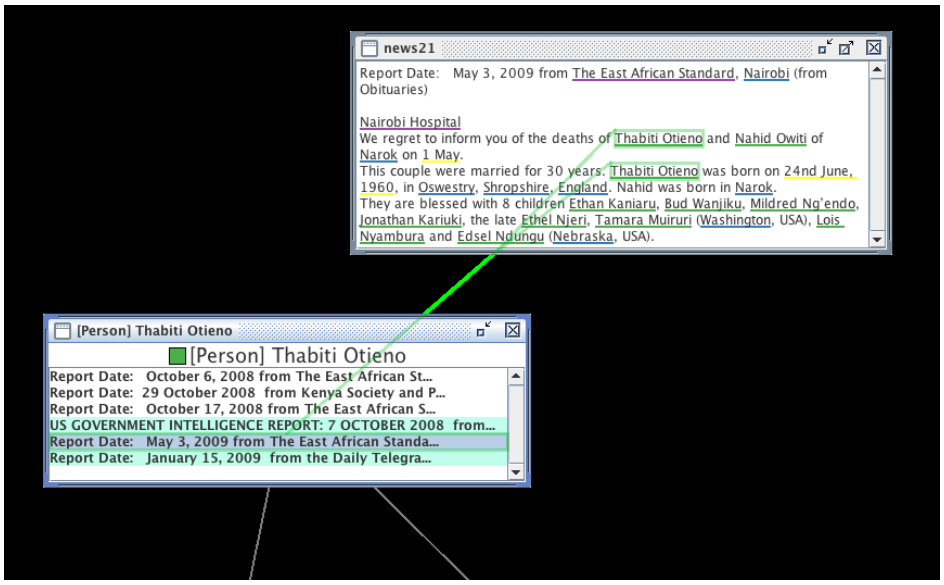

Figure 5.3: An entity showing internal links to instances in a document, as well as links to other open documents.

to show containment relationships (entities or search terms that occur in a document), and co-occurrence relationships (two entities that appear in at least one shared document). In addition, the analyst can add new connections to represent arbitrary relationships between documents.

Because of the potential for visual clutter through the proliferation of automatically drawn links, the automatic links are only shown if they connect to the currently selected object. These links can also be selected, which changes their color and redraws them on top of the other objects in the space, making them easier to trace.

As can be seen in Figure 5.3, the links are also multiscale. Similar to the cross application links described by Waldner et al. Waldner et al. [2010], selected links connect directly to the link target, not just to the frame that contains it. When a selection event occurs, the type of object being selected (document, entity, search term) is broadcast and the frames that contain the target return a list of the internal targets. Links are then drawn directly to all of those inner targets. These links direct the analyst's attention right to the actual occurrences of the entity or search term within the matching documents.

5.3 Details and Design Principles

In this section, we will discuss some of the specific design choices for AW, how they were addressed in the design of the tool, and how they support the goal of integrated foraging and synthesis. We will also discuss how the use of a large, high-resolution display as the foundation of the tool impacted many aspects of the overall design choices made in developing AW.

5.3.1 Spatial environment

The primary design decision that shaped every aspect of Analyst's Workspace was the choice to use a large, high-resolution display. This decision, of course, is the direct consequence of the earlier studies. We have discussed a number of advantages of using space as a sensemaking medium, including the role of human perception and the opportunity to perform incremental formalism. With respect to the primary goal of AW, space lowers the barrier to externalizing aspects of synthesis both cognitively and pragmatically. On the cognitive side, the ability to express lightweight relationships without the burden of knowing what the relationships is or how it fits into the rest of the investigation means that the analyst can externalize half-realized connections without interrupting the flow of an investigation.

On the pragmatic side, the ease with which these relationships can be expressed makes it more likely that the analyst will actually do it. This is a key design principle that underlies much of the design of AW – we encourage externalization of internal synthesis by reducing the amount of effort required to do so. In this case, we consider this effort to be both cognitive and physical (in the form of distracting interactions). While soliciting more information about a relationship might increase AW's ability to provide additional assistance, AW, like Entity

Workspace, purposely errs in the other direction in order to keep the cost of externalizing down [Bier et al. 2006]. By encouraging the analyst to externalize schemata into the space, AW helps analysts to represent relationships and categorize information as it is encountered, making relationships visible, and thus more accessible for analysis and questioning [Heuer 1999]. At the same time, it should help the analyst to abstract the information, making it easier to internalize.

Creating a spatial environment

As previously discussed in Chapter 2, a primary consideration in creating a spatial environment is the maintenance of a small figure to ground ratio. While a small ratio implies that we can display more objects, it is perhaps even more important that it creates more potential locations for those objects. Whitespace (or ground) is what allows one cluster of objects to be distinct from a second cluster of objects, and thus provides the expressive power of the space (i.e., the compliment of the Gestalt principle of proximity). A key advantage of using the large, high-resolution display is that it allows us to maintain a sense of space while at the same time using quite large objects – in this case, full text documents.

To maintain the ratio while retaining the visibility of the text, we made a couple of considerations. First, it is important to avoid extra decoration around the documents, as every extra pixel would be duplicated for every displayed document. For example, every document list can show a preview of a document before it is opened. This is implemented as a popup window rather than an additional frame so that it is not visible when it is not needed. Second, the documents can all be resized in conventional ways, but we also added a "smart zoom" feature that resizes frames to just fit their contents. While the feature is provided for convenience and to reduce wasted screen real estate, it also has the effect of making documents more recognizable since they are not all the same size.

Font size is another consideration. We used a font that would be considered comfortable for reading on a conventional display. The tradeoff of which is that if one steps back to see the entire display, the text may be too small to read. It is our experience that when many users first approach a large, high-resolution display, their inclination is to try to find a location from where they can "see it all", which, because of the limits of visual acuity, effectively *reduces* the amount of information that can be perceived. By using a smaller font we are encouraging users to sit closer and work with the display as an environment, where physical navigation allows them the make full use of the available resolution.

Supporting spatial organization

While AW imposes no spatial strategies on the analyst, it has been designed to specifically address two types of spatial structure that we most frequently observed in our previous studies: clusters and timelines. In this section, we will discuss some of the primary ways AW supports the development of these structures.

On top of the basic ability to move documents around the workspace, AW adds a couple of simple capabilities to make clusters easier to work with and use. Unlike the conventional window manager, AW allows all objects to be multi-selectable so they can be moved in groups like icons on the desktop. While this is a simple interaction, it is also quite important. First, it makes it much easier to restructure clusters as new information becomes available. Second, it allows whole clusters to be moved about the space as a single unit. This encourages regard of the cluster as the abstraction that it represents and increases the expressiveness of the workspace as it becomes feasible to express relationships between clusters. For example, documents could be clustered based on the principle suspects, and the clusters arranged to reflect the relationships between them.

AW also provides the standard alignment and distribution layout operations. These help the analyst to quickly arrange the documents in a cluster. While these operations are ostensibly about creating cleaner looking clusters, this is really an application of the Gestalt principles of similarity and closure to create more visually salient clusters.

In addition to these cluster-oriented operations, AW also provides a timeline operation that will reorder any documents in the currently selected set of objects into temporal order (undated documents and other objects will just be left in place). AW uses the bounding box defined by the selected documents to determine orientation of the timeline, so horizontal or vertical timelines can be created based simply on the original distribution of the documents.

It is important to note that all organizational operations are restricted to work only over the currently selected set of documents, similar to the spatial operations provided in Dynapad [Bauer et al. 2005]. In other words, these are analytic aids that are built into the space, but they are not intended to supplant the meaning (or lack thereof) of the space. This allows AW to support multiple spatial metaphors, including hierarchies of spatial metaphors in the space, another critical feature of space for supporting incremental formalism and the analytic process. For example, a cluster based on a person could be internally ordered temporally, highlighting the sequence of the person's actions, making it easier for the analyst to shape a hypothesized narrative.

5.3.2 Persistence

Persistence is another key design goal that motivates the use of the large, high-resolution display and the spatial environment. By persistence, we mean the continued availability of information within its semantic and spatial context.

The primary form of persistence is afforded by the use of the large, high-resolution display. Objects that are placed at a location on the display have a real physical location in the environment surrounding the analyst. In Analyst's Workspace, it is not just the document's position that remains persistent, but also its representation, in that the full text of the documents remains visible along with all of the annotations made by the user. This data persistence encourages the reading and rereading of the documents, which is a key part of understanding a textual data set [Stasko et al. 2007]. It is also worth noting that it can sometimes be difficult to initially know what piece of information in a document will prove to be the most important until some later document provides some context, making it important to keep details available.

Another important form of persistence supported by AW is what we might call investigative or process persistence. In other words, investigation state can be recovered directly from the workspace. The main mechanism for this is the treatment of entities and search results as persistent objects in the workspace. This makes the cognitive cost of branches in the investigation fairly low. At any time, the analyst can return to an earlier search and continue to pursue it, guided by the color-coding of the results.

Of course, the advantage of the persistent result sets extends beyond the ability to resume a previous branch of the investigation. The entities and the search terms can become a part of the sensemaking process, serving as representations that highlight key points or serve as bridges between spatial structures.

5.3.3 Contextual visualization

Another of the major design choices that we made to address this goal was to employ contextual or embedded visualization. Yost has previously demonstrated that the increased

screen real estate of large, high-resolution displays provides a distinct advantage for embedded visualizations (what she called "space-centric" visualizations) [Yost et al. 2007]. There are a number of reasons why embedded visualization works particularly well on large, high-resolution displays, but we will focus primarily on context. The available room allows us to embed additional content where it is important, where the analyst can see it and take advantage of it while conducting the investigation.

Contextual visualization occurs in many places in AW, and at many levels. It also does not always look like traditional visualization. At the highest level, the entire workspace is an example of contextual visualization. Taken as a whole, the spatial relationships between objects in the workspace are a free-form visualization of both the analytic process to date and the analyst's understanding of how the documents relate. Switching scale, the analyst has access to the textual contents of the documents, and within that, the extracted entities and any user highlights.

We also use contextual visualization to provide the analyst with additional information during the investigation. For example, the visible links between the documents and the entities and search results form a graph that underlies the sensemaking space, allowing the analyst to see connections not represented in the space.

We observe another form of contextual visualization in the use of mixed spatial metaphors in the space. In this case, it is the analyst who creates the visualization, facilitated by the tool. In Figure 5.4, we see a cluster based on a suspicious individual. Internally, the cluster has been ordered temporally into a vertical timeline. The analyst has then offset documents horizontally within the timeline to visually highlight separate sequences of events, forming an externalization of the analyst's assessment of causality and connectedness.

The key point about these different levels and forms of contextual visualization is that AW is using the same visual artifacts for all of them. In other words, the same object

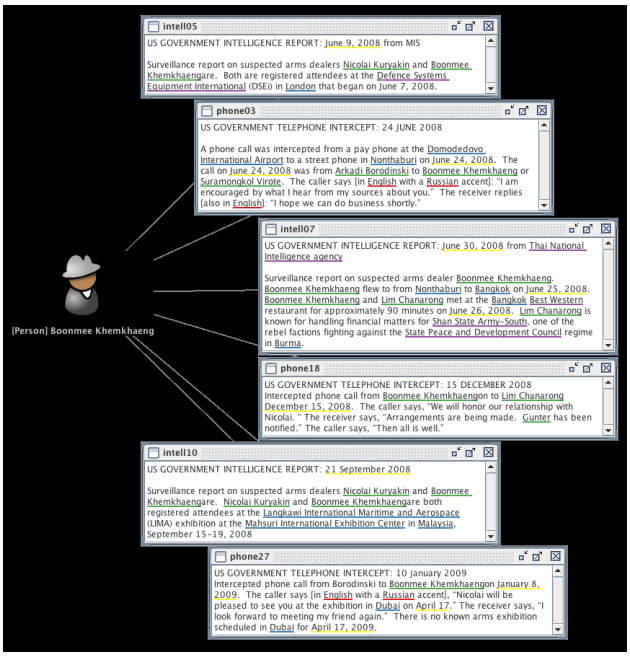

Figure 5.4: Timeline using offsets to mark different threads. Note also the iconified entity labeling the structure.

is simultaneously a representation for some underlying concept, part of a larger semantic structure developed by the analyst, a piece of data, and a tool for further exploration. It is this feature of the workspace that most strongly supports the close integration between the foraging and synthesis aspects of the sensemaking process. Since the artifacts are the same, there is no enforced break between actions relating to synthesis and actions relating to foraging - no need for the analyst to break out of the context of the investigation to schematize information, and vice versa, no need to access a separate tool to find more evidence in support of a hypothesis.

5.3.4 Think global, act local

We have described a number of key ways that Analyst's Workspace leverages the large, high-resolution display to provide integration and context. In this section we want to explicitly address some of the more pragmatic issues of designing for these displays. In particular, we

want to focus on the use of the display for a single workspace or view. The rationale for why we use a single workspace over a collection of linked views rests primarily in our desire for integration and our analysis of the expressive qualities of space.

While we have provided the user with a single workspace that we expect to be regarded as a unified entity, our emphasis on physical navigation means that our analyst will primarily be working close enough to the display that most of it will be out of his or her field of view. Because of this, there are a number of design considerations that we must address.

The visual links that are used to connect various objects across the space are an example of directing the analyst's attention across the display to related information that was outside of his or her field of view. This also encourages the analyst to spread out to use the display more effectively, since losing track of a particular piece of information is less likely.

On a strictly pragmatic level, we also find that most of the conventional assumptions about controls and notifications must change as well. It is fairly standard to place controls around the periphery and notifications in the center of the display. These assumptions are no longer valid, and can result in loss of efficiency and missed notifications [Sabri et al. 2007]. In AW, we address this in two ways. First, all controls are available locally, either through contextual menus or hot keys. Second, the system uses a combination of the currently focused object and the cursor location to derive the user's current focus of attention. Notifications, dialogs, and new frames are all placed near this point, which reduces missed messages and makes interaction much more efficient.

The dynamic between global information and local interaction also relates back to the previous discussions of operations that are performed only over the selected set of objects. While the space frequently takes on some form of global semantics, the attention of the analyst at any given moment is typically a smaller set of documents or objects, so it makes sense for operations to only operate over that current working set.

5.4 Analyst's Workspace in Action

The development of Analyst's Workspace was driven by extensive internal use of the tool as we worked to understand the opportunities offered by large, high-resolution displays for sensemaking. Most significantly, it has been used to successfully solve the textual analytics problems in the past two VAST Challenges [Costello et al. 2009], as well as being incorporated into a side project investigating the integration of storytelling algorithms into the environment. In this section, we will describe our work integrating the storytelling algorithms, our experiences with the most recent VAST challenge and discuss a longitudinal evaluation of the environment by an outside user that we recently conducted.

5.4.1 Storytelling

The storytelling project expands on the theme of creating an environment that works with the analyst by adding storytelling capabilities to AW [Hossain et al. to appear]. The concept was to find ways to incorporate the storytelling algorithms into the interface so that they would support a form of analytic discourse with the analyst, suggesting connections when requested by the analyst.

There were several additional capabilities built into AW in the course of this collaboration. The core storytelling algorithm was designed to find connecting documents between any two documents selected by the analyst. The analyst simply draws a link between the two documents and AW populates the line with any connecting documents. Selecting any of the connecting links highlights the terms shared by the connecting documents that were used to create the link, assisting the analyst in assessing the value of the connection. If the analyst does not approve of the link, the document can be removed, or the link severed. The core

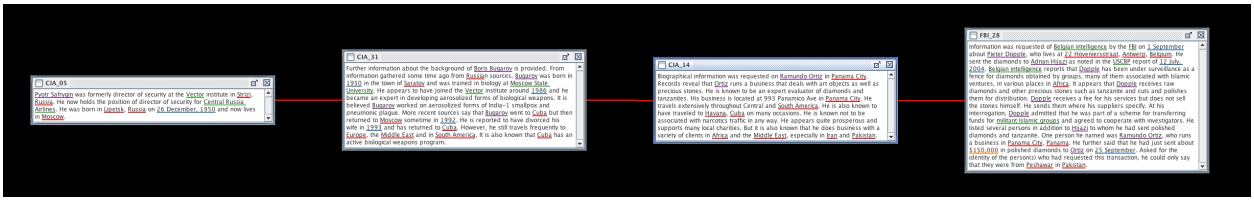

(a) The analyst has requested a connection between documents CIA_05 and FBI_28.

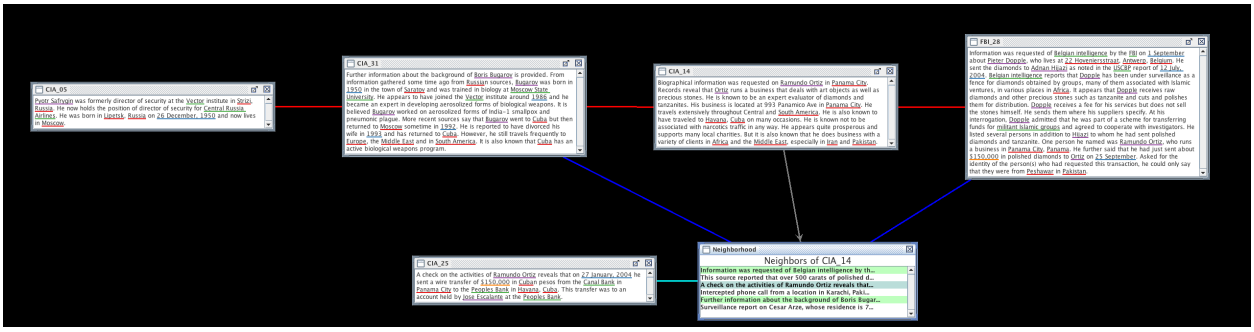

(b) The analyst thinks the connection between CIA_31 and CIA_14 is weak. Examining the neighborhood of CIA_14, the analyst discovers an interesting transaction in CIA_25.

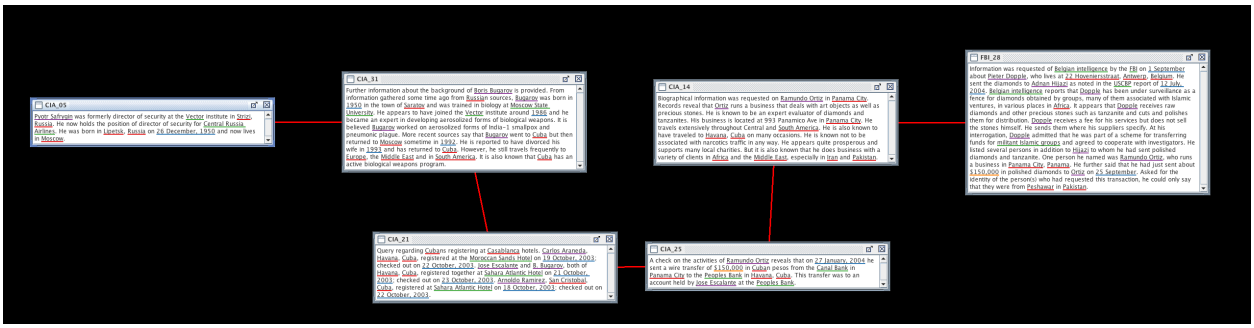

(c) Following the money, the analyst finds a stronger connection back to CIA_31 and edits the links to show this new story.

Figure 5.5: A sequence of screenshots showing an analytic discourse between the analyst and the storytelling alogorithm.

idea is to never take control away from the analyst and to make the working of the algorithm as transparent as possible (Figure 5.5).

We also integrated a tool for identifying "neighboring" documents, which finds documents that are similar to the current one based on an analysis of word frequencies. This tool was originally designed to help the analyst to tune the output of the storytelling tool by exploring the neighborhood of documents that did not fit, or connections that seemed weak. However, this turned out to be a useful tool in and of itself for finding potentially interesting

documents given a document that the analyst had identified as important.

In addition to these algorithms, which allow the analyst to navigate the dataset from document to document, we also added some tools to traverse the entity space, making more concrete connections through the dataset. Unlike the storytelling algorithms, these tools were relatively straightforward graph traversal tools based on the co-occurrence of entities. Given two entities (identified, again by drawing a link between them), the algorithm would return a sequence of entities that connected the two. The more precise version of the tool is based on a semantic parse of the documents, and only makes connections based on co-occurrence within the same sentence. As this fairly strict interpretation frequently fails to make a connection, the tool falls back on document-level co-occurrence.

Our work with these interfaces is currently incomplete. We found that there is a disconnect between algorithmic functionality and capabilities and analytic needs and expectations. Sometimes, as shown in Figure 5.5, the algorithmic support provides an interesting connection and directs the analyst's explorations. However, when the algorithm suggests possible connections that are completely irrelevant, it is less clear how to support the analyst in redirecting the results without distracting from the core investigation. There is a delicate balance between the need to leave control in the hands of the analyst and not requiring excessive interaction, and this project is still grappling with this question. However, this experience does demonstrate the use of AW as a testbed for developing new interactions and tools with the spatial context provided by the large, high-resolution display.

5.4.2 VAST challenge 2011

The text analysis component of the VAST Challenge 2011 required contestants to identify potential threats to the city of Vastopolis buried within 4,474 news articles. Using AW, we

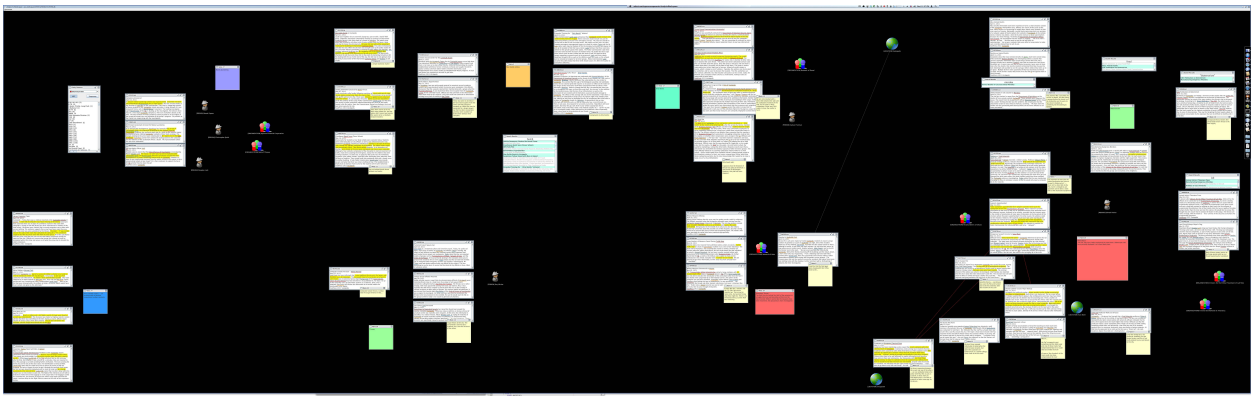

Figure 5.6: Our solution to the text analysis component of the VAST 2011 contest.

were able to correctly identify the threat to the city, and we were selected for an award: "Novel Use of Large Screen Workspace to Support Analysis" [Andrews et al. to appear].

The most significant challenge when faced with a dataset containing this many documents is finding a place to start. To this end, we used the AlchemyAPI [alc 2011] to classify the documents into basic news categories, and added a simple browser to AW to view documents by category. Our investigation then began with the "law and crime" documents, which effectively cut the number of documents we needed to look at from 4,474 down to 81. Scanning through these resulted in approximately five leads about threats to Vastopolis. Each lead was placed in a different region of the display based on the implied threat. After the initial pass through the category, we visited each grouping in turn, spidering out from each document based on interesting entities or keywords.

The most significant lead came from the arrest of some members of a local terrorist group called the Paramurders of Chaos (PoC), who were caught with the remains of some high technique laboratory equipment. Following the PoC entity, we find they are currently under suspicion and one of their members was caught trespassing at a local food plant. This suggests the possibility of bioterrorism, and this leads to a local professor who is an expert in the field who recently had his lab burgled. Ordering these document temporally in AW

yields a plausible timeline of events, from robbery, to laboratory, to the trespassing incident. Interestingly, this dataset is particularly weak on entity-level connections, and there were minimal leads to pursue from these documents. While pursuing other lines of investigation, however, we uncovered other evidence, such as the mass die off of fish, that we could tie back to this plot.

Figure 5.6 provides a screenshot of the final state of the workspace. The primary threat can be found on the right of the space. The rest of the space is arranged so that the threat level decreases as one moves to the left. Immediately next to the bioterrorism evidence, we have a threat related to weapons smuggling, and above that, a dirty bomb threat. Both of these have already been resolved by the authorities. Moving further left, we have some other bomb incidents, some money laundering, and on the far left, a collection of isolated criminal activities.

There are a couple of key points to make about this workspace. First, note the layers of meaning. At the top level, the entire space is roughly reflective of the severity of the threat. Similar incidents are loosely grouped, so, for instance, all of the bomb incidents are near one another along the top. Internally, documents are arranged in columns, which reflect the temporal ordering of the events or reports. Horizontal displacement or vertical separation within these columns are used to show disconnects or incidents that are not directly part of the same sequence of events, but are still related. The second thing to notice is that use of the iconified entities, which are being use to label the structures and show connections, showing the transformation from foraging tool (list of documents) to synthesis artifact. The final point is that because of the available room, we are able to keep the whole of the investigation visible, even if it seems resolved or disconnected from the primary threat. This means that if new information is uncovered (as it was in our primary plot), it can be added directly to the corresponding collection of events in the space. In addition, it provides a

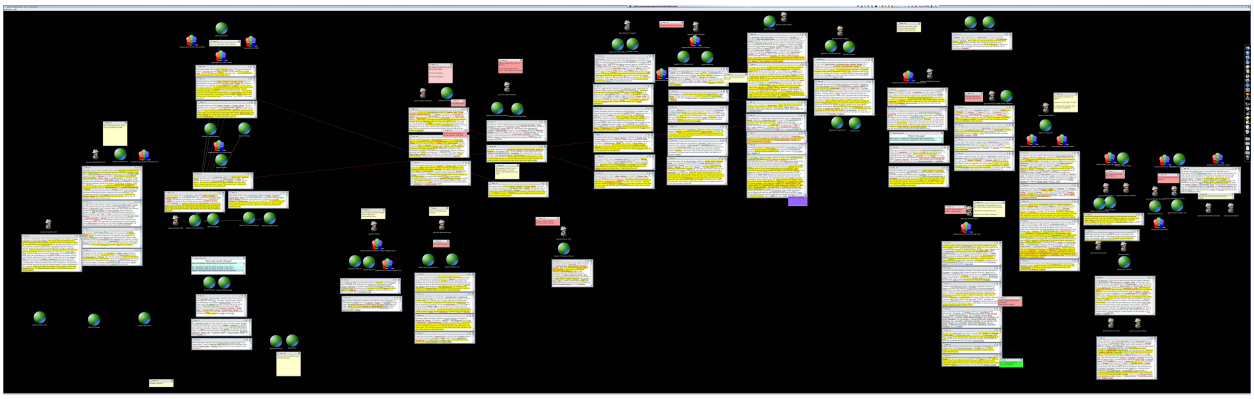

Figure 5.7: Final solution provided by longitudinal subject. Note the distinct columns, each marked by a collection of related entities.

record of the entire investigation, so other analysts can see not just our solution, but our entire process, including what we rejected and why.

5.4.3 Evaluation

Conducting a controlled evaluation of a tool like AW presents a number of significant challenges. As our focus is on supporting an integration of foraging and synthesis, the problem to be solved needs to be complex enough to make internal synthesis impractical, implying a significant time commitment from an evaluation subject. In addition, as AW does not impose any structure on the space, it is left to the user to develop an effective strategy on how the space should be used, which requires enough insight into the analytic process to know what may be important to capture.

Rather than attempting a full comparative evaluation, we focused on performing a longitudinal exploration aimed at examining the effectiveness of the support for integrated foraging and synthesis. Due to the time involved, we conducted this evaluation with a single subject, across nine sessions, each lasting from one to two hours. No time limits were imposed on the subject, she was just asked to be a thorough as possible in developing her case. The subject

began with a training exercise, consisting of 41 synthetic intelligence reports (28 of which were important). The training spanned four sessions, and 4.5 hours. An extensive debriefing and a discussion of the effectiveness of the various strategies she employed in the course of the investigation followed this. She then moved to the primary investigation, consisting of 111 documents (57 important), spanning five sessions for a total of 9 hours.

Her investigation began by opening the three person entities with the highest frequencies. From there, she followed the FCFT strategy - opening connecting documents and using those to find new entities of interest. As she worked, she followed a general schematizing strategy of grouping documents into temporally ordered columns based on the primary person. She placed columns of related people together, so that the plots and events she was tracking divided the space up into distinct regions. Locations were another important component of her investigation, and when a person was particularly associated with an address or other location, she opened the associated entity and placed it beside the person entity.

Throughout the investigation, she made heavy use of the visual links. Whenever she opened a new document, she would look at all of the links to open entities to see how the document fit into the overall picture she was forming. She also made heavy use of the links while she was sorting out relationships – using the deep linking to rapidly see all of the activity surrounding a particular entity.

The final state of her investigation can be seen in Figure 5.7. At this point, she has positioned the columns so that the more important people are closer to the top of the workspace, and the horizontal placement shows rough relationships between the displayed people. Her solution is not perfect; she is a novice to intelligence analysis and she made some incorrect connections based on false assumptions that encouraged her to try and weave some additional events into her narrative, clouding the final form of the coming attack. However, her solution did

have the network of the players and most of the sequence of events leading up to the attack. The solution contained 56 of the 57 important documents and about 15 "noise" documents.

More importantly, her process showed clear signs of the tight integration between foraging and synthesis. Her process was a constant flow between digging into document, tracing entity connections, and schematizing the results. Even as she was organizing the workspace into its final form, she was continuing to use the entities to flesh out connections, and in one notable instance elevating the importance of a person as she uncovered additional evidence of just how connected he was. While this study provides validation that AW can be successfully employed to conduct investigations, future development would certainly be enhanced by evaluation with additional, and preferably more experienced, users.

5.5 Future Work

There is also considerable work that can be done, both conceptually and developmentally, to extend this work. We have described AW as integrating the foraging and synthesis stages of the sensemaking process. However, if we were to place it on the spectrum of tools that we described earlier, we would describe it as occupying the middle of the process. It serves as a bridge, connecting the foraging and the synthesis sub-processes, but it is perhaps weakest at the two ends of the process.

This is most obvious at the start of an investigation. If there is no "seed" document or directed task, the analyst has few options for finding a starting clue. The most promising place to start is by examining entity frequencies in the entity browser. This is generally fairly effective, but can lead to some false starts if the data set is noisy. Another approach would be to just start reading documents, but this approach tends to be fairly ineffective

and time consuming [Kang et al. 2009]. Ideally, we would want to find some way to provide the analyst with an overview of the data set. The challenge would be to integrate it into the overall process we have already developed, rather than being a tacked on overview of the data set that provides the initial lead and then is not used again.

On the other end of the process, we have not addressed the problem of transforming the investigation workspace into a presentable report. There are two related issues to be explored here. The first question to explore is the question of how personal the semantics of the spatial layouts are. In other words, are final workspaces meaningless to everyone but the analyst who created it, or can others "read" the solution out of the space? Exploring this will tell us how directly we can use the workspace as a result and how much the tool could potentially derive from it.

The second question is how much more information can we reasonably get from the analyst. As we previously stated, our focus was primarily on the use of the space, and we were purposely conservative in how much information we required the analyst to provide. An interesting avenue of exploration would be to find additional minimally invasive ways to elicit additional information from the analysts. Of course, the benefits of this would extend beyond support for transforming the workspace into a report and could impact the entire sensemaking process.

The previously mentioned work on Sandbox and Entity Workspace provide some potentially interesting approaches to this problem. Another promising direction is to make use of the user's spatial manipulation as an input back into the system. While some work has been done with spatial parsers to infer structures based on placement [Francisco-Revilla and Shipman 2005], we are currently working on a more directed approach that uses spatial manipulation directly as an input to a storytelling algorithm. In other words, the goal is to support an

analytic discourse with the analyst, rather than a post hoc summarization of the final state of the workspace.

5.6 Conclusion

Analyst's Workspace is an analytic tool that leverages the spatial environment provided by a large, high-resolution display to create a sensemaking workspace. The primary objective of the tool is to provide a fluid working environment that tightly integrates both foraging and synthesis, and encourages the externalization of artifacts of the synthesis process.

While there remains more work to be done in evaluating AW's effectiveness, we feel this work already has a contribution to make. Our perspective on the importance of supporting the whole of the sensemaking process is one that has been overlooked and we hope that our discussion of the design goals of AW can help to shape future tool development. In particular, we would like to highlight a couple of key elements in the design that we feel are particularly important:

- The low interaction overhead for expressing relationships encourages externalization.
- The flexibility of the spatial environment encourages the adoption of incremental formalism and exploration.
- The use of the same visual artifacts for synthesis and foraging leads to more fluid movement between the two processes.
- The large, high-resolution display creates the opportunity for layers of context through the availability of high levels of detail (e.g., raw text, highlights, entities, documents, spatial structures).

While all of these points are important, it is this last point that provides the real strength of large, high-resolution displays in achieving the others. It is through the multi-scale representations that are supported by the display that we can create artifacts that are usable for both foraging and synthesis. This in turn is what allows us to make full use of the environment for low cost externalization.

In addition to emphasizing a different perspective on the goals of analytic tools, we hope that our description of Analyst's Workspace has made an argument for the use of large, high-resolution displays as a platform for the development of analytic tools. We believe that AW demonstrates that these displays are not just slightly larger monitors, or purely collaboration aids, but unique environments with a new set of affordances, opportunities, and design challenges, and we hope that this work serves to inform and inspire future use of them.

Chapter 6

Conclusion

This work was driven by three research questions:

- What cognitive affordances are provided by large, high-resolution displays and how can they change the way users work and think?
- What advantages, if any, do large, high-resolution displays provide for the sensemaking process?
- How can we design analytic tools that leverage these advantages?

Through the studies and the development of Analyst's Workspace, we have strived to answer these questions. In this chapter, we will summarize the primary contributions of this work and discuss how they address these questions. We will then conclude with some future areas of exploration suggested by this work.

6.1 Contributions

Spatial environment

We have demonstrated through observational studies and analysis of the literature that large, high-resolution displays provide an environment that is implicitly spatial. The perception of space comes from two primary sources. The first source comes from the use of representations that are dwarfed by the available display space. Due to the high pixel density of the display, common representations consume very little of the display. As a result, position on the display assumes meaning, and meaningful spatial relationships can be conveyed. Second, the physical properties of the display subsume the visual field, and require the user to perform physical navigation, thus engaging a number of embodied resources developed for navigating the spatial environment of the everyday world. It is this aspect of the display that provides new cognitive affordances and thus creates the most potential for changing how the user engages with the displayed information, and provides an answer to the first of the research questions.

Space for sensemaking

As noted in Chapter 2, the use of space to support sensemaking, and in particular, the process of synthesis, is not a new one. However, this work is the first to explicitly combine these two for use on a large, high-resolution display, thus addressing the second question. Within the context of the large display, this work also separates out two uses of the space for sensemaking: memory, and organization. The use of space for external memory is a fairly obvious one. However, the large display offers a couple of key advantages. The large number of pixels permits a large amount of information to be displayed simultaneously, where it

can be retrieved using physical navigation driven by spatial indexing and the surrounding spatial context. The advantages of spatial organization relate partly to the ease with which humans can perceive spatial patterns and relationships and partly to the lack of explicit structure imposed by the space. This allows the user to employ multiple spatial metaphors and engage in incremental formalism, developing meaningful structures as understanding of the underlying data grows.

Importance of physicality

A key contribution of this work is to demonstrate that there are critical differences between the physical space provided by the large, high-resolution display and virtual workspaces that are used to provide the appearance of space on conventional displays. The study described in Chapter 4 clearly demonstrated that there were physical and cognitive changes brought about by the different forms of space, leading to different user behaviors. This led to a more coherent view of the overall space, and more effective externalization during the sensemaking task. This difference is critical for a number of reasons. First, it highlights how support for a greater range of embodied resources strengthens the coupling between the user and the environment. Second, it suggests that tool developers making use of the virtual workspace metaphor should be cautious about claims or design decisions that are based on human spatial abilities. Finally, this clearly demonstrates that large, high-resolution displays are fundamentally different from conventional displays and tools for the two environments require different sets of design goals and principles.

The previous contribution only partially answered the second research question, because while it was true that the spatial environment created by the display could be leveraged for a spatial approach to sensemaking, our review of the literature indicates that this is not

unique to the display environment. However, the results of the study in Chapter 4 showed that the display supports spatial sensemaking at a different level than conventional displays, allowing the user to interact with more detailed information, and to externalize aspects of the sensemaking process more efficiently.

Analyst's Workspace

The final major contribution of this work is the development of Analyst's Workspace. AW is a fully functional analytic environment, designed specifically for the large, high-resolution display, which supports an integrated approach to sensemaking, weaving together the activities associated with foraging and synthesis. The development of the environment also served as a vehicle to explore some of the design space for developing large, high-resolution display based analytic tools. To address the third research question, the key design principles that the work identified are:

- Make use of the spatial affordances of the display.
- Low cost interactions for expressing relationships can encourage externalization, supporting synthesis.
- Use the large number of pixels to keep detailed information available through physical navigation.
- Investigative or process history is another valuable use of the display space.
- Encourage incremental formalism by not prematurely imposing structure or meaning to the space and by supporting localized operations, rather than global structuring tools.

- The availability of high levels of detail supports the creation of layers of context and integrated visualization, creating a multiscale environment that can convey large amounts of information at varying levels without compromising on the flexibility of the space.
- Detailed representations with integrated functionality provide the opportunity to use the same artifact for both foraging and synthesis, leading to a more fluid process.
- Single function or concept representations are more easily abstracted and thus more amenable to being integrated into semantic structures.

6.2 Limitations and Future Directions

Throughout this document we have identified some of the caveats and limitations of this work. In this section, we will attempt to summarize these as well as discussing some of the other potential avenues for future exploration.

One of the most obvious limitations of this work lies with the difficulty of evaluating the sensemaking process. Sensemaking is a complex cognitive process, and professional analysts are highly trained to do their jobs. Accessed to trained practitioners can be difficult to achieve, and they rarely have time for learning new systems or participating in lengthy evaluations. As we observed in our studies with novices, individual differences in analytic ability easily overwhelm any effect that we are trying to study, yielding very high variances in actual performance. A further challenge is that the range of cognitive activities involved in the sensemaking process and the dense interconnectedness between them lead us to question the validity of trying to decompose the process into subprocesses that may be more amenable to study.

In this work, that effect was felt most profoundly in the study of the effect of virtual versus physical navigation described in Chapter 4. In that study, we could show behavioral

differences, but our argument is weakened by the fact that we could not show a corresponding difference in analytic performance, which is inevitably what really matters. While the demonstration that the use of physical space led to behavioral changes and an increase in the amount of information being externalized, more work remains to be done to demonstrate that these changes lead to more effective sensemaking. One way to address this would be to find a more experienced subject pool, and we have discussed working with students who have been studying analysis techniques as one way to form a more balanced pool.

Turning to our work on Analyst's Workspace, there are a number of open areas of inquiry. Again, evaluation remains an open question. While we have shown some examples of AW in use, more thorough evaluation remains to be done to back up our numerous claims about the utility of the various features of the environment. Evaluating AW will prove particularly challenging, partially due to the issues mentioned above, but also to the nature of the environment. AW takes a somewhat different approach to the investigative process than other tools, requiring a second level of training of any potential subjects, thus requiring more time from the subjects.

Another open challenge is the fact that AW still does not cover the entirety of the sensemaking process. AW's most glaring lack is the impoverished facilities for assisting an analyst in starting an investigation. AW is basically a blank slate at the beginning of an investigation. While this works for focused investigations that have a direction or a starting point, it is obvious from our experiences that some facility for gaining an overview of the available data and being able to filter down to some interesting leads would be beneficial (i.e., support for top-down as well as bottom-up investigations). A related, but less critical deficiency can be found at the other end of the sensemaking process. AW supports incremental formalism up to a point. However, it leaves much of the final structuring in the hands of the analyst.

There is certainly room for some more advanced structuring tools that could help the analyst in the development of hypotheses and, later in the process, generating reports.

Examining the entire approach taken by AW, one of the more basic questions that still needs to be addressed is the question of scalability. With the focus on detailed data and the physical nature of the space, it is obvious that there are hard bounds on how much information can really be managed using the approach taken with AW. Even if more screens could be added to the system, there are human physical limitations to consider. The earlier use of the term "human scale" does imply an upper bound on the size of the environment. There are also cognitive limitations to consider. While representing layers of information through spatial relationships and contextual visualization helps the analyst to externalize aspects of the sensemaking process, it has replaced it with a need to manage a visual field full of information and a spatial environment that has to be managed.

This suggests a couple of lines of inquiry. One might ask at what point does the cognitive cost for keeping track of information in the space become too great? Another direction would be to examine the effect of aggregated representations and other "space saving" techniques on this process. Jakobsen certainly provides evidence that simply porting existing visualization techniques to the large display environment is likely to be problematic [Rønne Jakobsen and Hornbæk 2011]. Another interesting question that should be asked is how much information *needs* to be displayed. There seems to be an assumption that analytic tools should be able to scale up to thousands or even millions of documents because of the nature of the information space. It seems worth questioning this assumption. There is a difference between making sense of one million documents and having one million documents from which to find evidence. Do human capabilities or just the nature of analysis tasks produce smaller "working sets" of documents more amenable to the spatial approach?

A related line of exploration would be into the nature of representation. This work has demonstrated the utility of fully detailed representations. However, there are many more questions that could be asked. For example, how much information is "enough" in a representation? How much do the full text documents serve as reminders that are seen but not read, and how much are they actually consulted for detailed information? Is there a way to break out separate parts from a document and use them in different sensemaking structures while maintaining the context of the original document?

Finally, there is a question of how applicable this work is beyond the domain of intelligence analysis. The actual sensemaking activities that were observed in this work are primarily identifying important information, categorizing, and arranging; tasks that are fairly common for knowledge workers across many domains. However, it is worth investigating what it would take to support these different domains. Certainly, AW's entity-centric interaction would not be useful in all instances. There are also domains, such as cyber security and financial analysis that deal with much more structured data, and more work would need to be done to determine how well the advantages or the associated design principles identified as part of this work actually hold in those domains.

6.3 Final Words

One of the primary goals of visual analytics is to develop tools and environments that that work with a human operator, forming a tightly coupled cognitive system marked by the marriage of computational power and cognitive resources. To support this goal, it is important that the environment supports a wide gamut of embodied resources to support the exchange of information between human and environment.

In this work, we have argued that large, high-resolution displays provide a fundamentally different environment than conventional displays, and that the physicality of the display supports a more embodied interaction between the user and the displayed information. We believe that this greater embodiment supports better communication of information and provides a richer environment for externalizing aspects of the sensemaking process. There is still much to be researched, and we hope that this work serves as inspiration for future investigations into large, high-resolution displays and the development of analytic environments.

Appendix A

Comparative Study Documents

VirginiaTech

Office of Research Compliance
Institutional Review Board
2000 Kraft Drive, Suite 2000 (0497)
Blacksburg, Virginia 24061
540/231-4991 Fax 540/231-0959
e-mail moored@vt.edu
www.irb.vt.edu

FWA00000572(expires 1/20/2010)
IRB # is IRB00000667

DATE: August 15, 2008

MEMORANDUM

TO: Christopher L. North
Christopher Andrews

Approval date: 8/15/2008
Continuing Review Due Date:7/31/2009
Expiration Date: 8/14/2009

FROM: David M. Moore

SUBJECT: **IRB Expedited Approval:** "Display Size Comparisons for Intelligence Analysis" , IRB # 08-464

This memo is regarding the above-mentioned protocol. The proposed research is eligible for expedited review according to the specifications authorized by 45 CFR 46.110 and 21 CFR 56.110. As Chair of the Virginia Tech Institutional Review Board, I have granted approval to the study for a period of 12 months, effective August 15, 2008.

As an investigator of human subjects, your responsibilities include the following:

1. Report promptly proposed changes in previously approved human subject research activities to the IRB, including changes to your study forms, procedures and investigators, regardless of how minor. The proposed changes must not be initiated without IRB review and approval, except where necessary to eliminate apparent immediate hazards to the subjects.
2. Report promptly to the IRB any injuries or other unanticipated or adverse events involving risks or harms to human research subjects or others.
3. Report promptly to the IRB of the study's closing (i.e., data collecting and data analysis complete at Virginia Tech). If the study is to continue past the expiration date (listed above), investigators must submit a request for continuing review prior to the continuing review due date (listed above). It is the researcher's responsibility to obtain re-approval from the IRB before the study's expiration date.
4. If re-approval is not obtained (unless the study has been reported to the IRB as closed) prior to the expiration date, all activities involving human subjects and data analysis must cease immediately, except where necessary to eliminate apparent immediate hazards to the subjects.

Important:
If you are conducting **federally funded non-exempt research**, please send the applicable OSP/grant proposal to the IRB office, once available. OSP funds may not be released until the IRB has compared and found consistent the proposal and related IRB applicaton.

cc: File

Informed Consent for Participant of Investigative Project
Virginia Polytechnic Institute and State University

Title of Project: **Display Size Comparisons for Intelligence Analysis**

Principal Investigators: Christopher Andrews and Chris North

I. THE PURPOSE OF THIS RESEARCH/PROJECT

You are invited to participate in a study of examining the impact of display size on analytic activities. The results of this study will inform the development of future analytic tools.

II. PROCEDURES

You will be asked to perform a basic analytic task typical of an intelligence analyst. The task will involve sitting at a desk for approximately four hours (or until you are satisfied that you have exhausted the possibilities of the document collection). At the end of the session, you will be "debriefed". You will need to report any hypotheses you have, what evidence supports these and your strategies for arriving at them.

The entire session will be video tapped, and you may occasionally be interrupted and asked about your progress.

III. RISKS

The proposed experiments involve normal office-like activities and, as such, pose minimal risks. The only foreseeable risk is a mild discomfort caused by the heat emanating from the large display. If you experience headaches or other discomforts, please feel free to step back from the display and inform the experimenter.

IV. BENEFITS OF THIS PROJECT

Your participation in this project will provide information that may be used to inform the design of tools for intelligence analysis on large displays. No guarantee of benefits has been made to encourage you to participate. You may receive a synopsis summarizing this research when completed. Please leave a self-addressed envelope with the experimenter and a copy of the results will be sent to you.

You are requested to refrain from discussing the evaluation with other people who might be in the candidate pool from which other participants might be drawn.

V. EXTENT OF ANONYMITY AND CONFIDENTIALITY

The results of this study will be kept strictly confidential. Your written consent is required for the researchers to release any data identified with you as an individual to anyone other than personnel working on the project. While a video record will be made, only subject numbers will be used in analysis and reports.

VI. COMPENSATION

Your participation is voluntary. Participants will be paid at the rate of $10/hour, for a total of $40 for completing the task. If you leave early, you will be paid the hourly amount, unless you have completed the task, in which case the full $40 will be paid.

VII. FREEDOM TO WITHDRAW

You are free to withdraw from this study at any time for any reason.

VIII. APPROVAL OF RESEARCH

This research has been approved, as required, by the Institutional Review Board for projects involving human subjects at Virginia Polytechnic Institute and State University, and by the Department of Computer Science.

IX. SUBJECT'S RESPONSIBILITIES AND PERMISSION

I voluntarily agree to participate in this study, and I know of no reason I cannot participate. I have read and understand the informed consent and conditions of this project. I have had all my questions answered. I hereby acknowledge the above and give my voluntary consent for participation in this project. If I participate, I may withdraw at any time without penalty. I agree to abide by the rules of this project

_____ _____

Signature Date

_____ _____

Name (please print) Contact: phone or address or

_____

email address (OPTIONAL)

Should I have any questions about this research or its conduct, I may contact:

Investigator: Christopher Andrews Phone (978) 314-7881
Graduate Student, Computer Science Department
email: cpa@cs.vt.edu

Review Board: David M. Moore, Office of Research Compliance, CVM Phase II (0442) 231-4991

cc: the participant, Christopher Andrews

Demographic Survey

Please help us to categorize our user population by completing the following items.

Subject ID _____

Gender (circle one): Male Female

Age: _____

Occupation (if student, indicate graduate or undergraduate):

_____

Major / Area of specialization (if student): _____

Have you taken CS 5764? _____

When? _____

Describe any other experience relevant to intelligence analysis.

_____

_____

Which operating systems are you familiar with? _____

Which do you use on a daily basis? _____

Do you use virtual desktops? _____

If so, how many workspaces do you use? _____

What kind of display do you typically use [laptop, LCD, CRT]? _____

Do you use more than one display? _____

For each display you have in your typical workspace, list the size and resolution [estimate if necessary].

_____

_____

_____

Describe any experience with large displays [gigapixel, storm, the table, etc.].

_____

_____

_____

Debriefing Questionnaire

You may use any notes that you have generated to help you to complete this form.

Subject ID _____

In as much detail as possible, describe your strategy for solving this scenario.

How did you find documents of interest?

How often did you reread of consult documents you had read previously?

What did you do with documents that seemed relevant?

What did you do with documents that seemed irrelevant?

Please describe any tools that you wish you could have had for performing this analysis.

In as much detail as possible, describe your hypotheses about what is happening in Alderwood. Provide evidence as necessary.

Appendix B

Comparative Study Screenshots

In this appendix we list all of the final screenshots from the comparative study. The only exception is the large display user L1. Due to a technical problem during the study, we were not able to retain his screenshot.

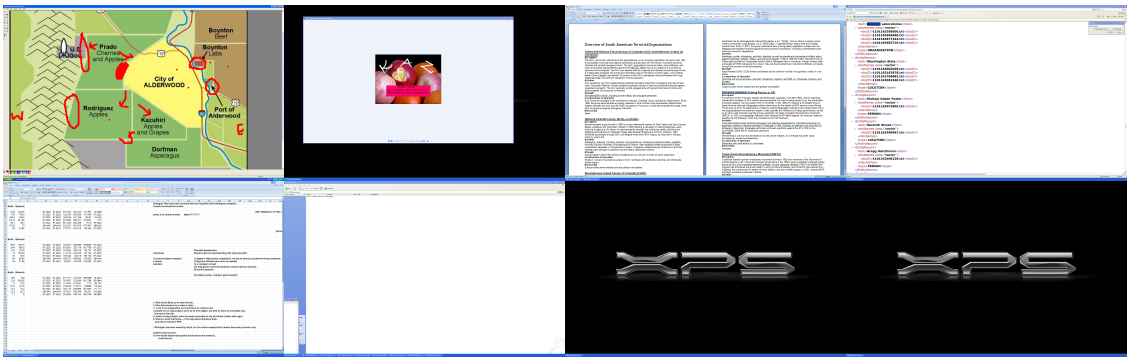

Figure B.1: Final state of the workspace for subject L2

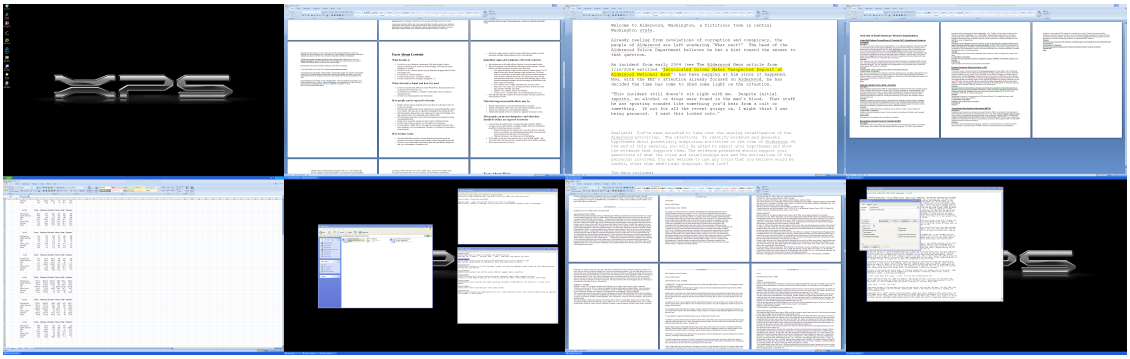

Figure B.2: Final state of the workspace for subject L3

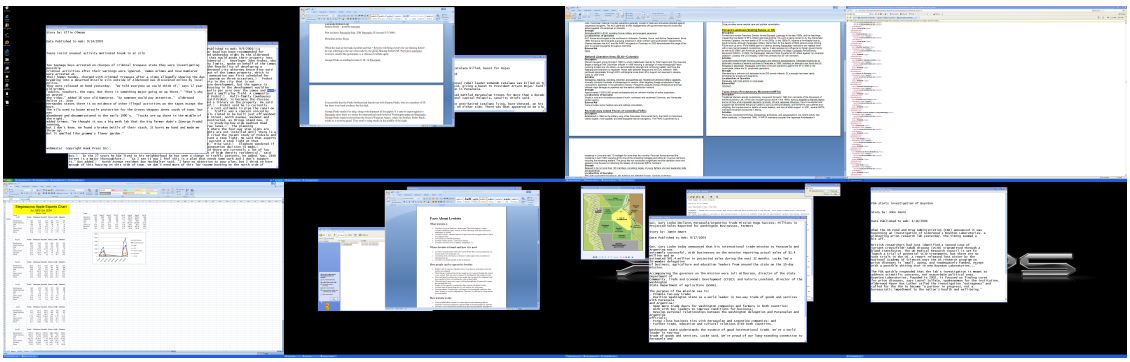

Figure B.3: Final state of the workspace for subject L4

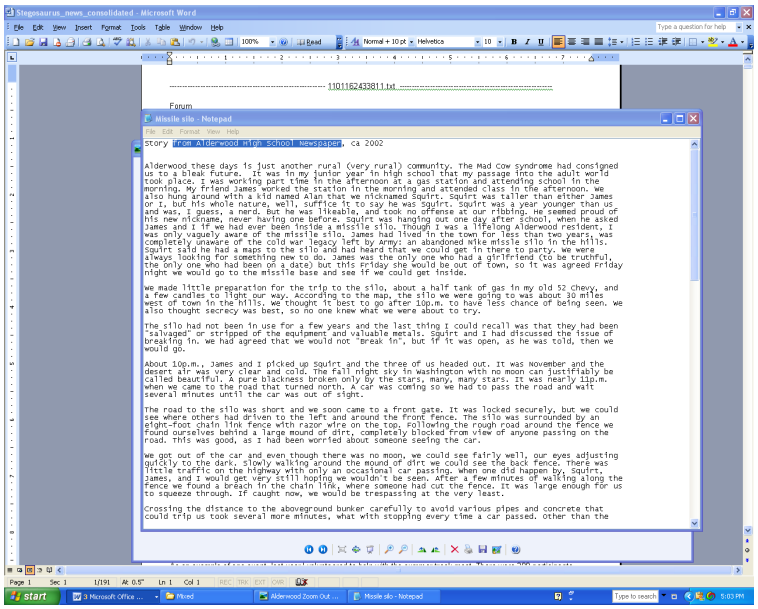

Figure B.4: Final state of the workspace for subject S1

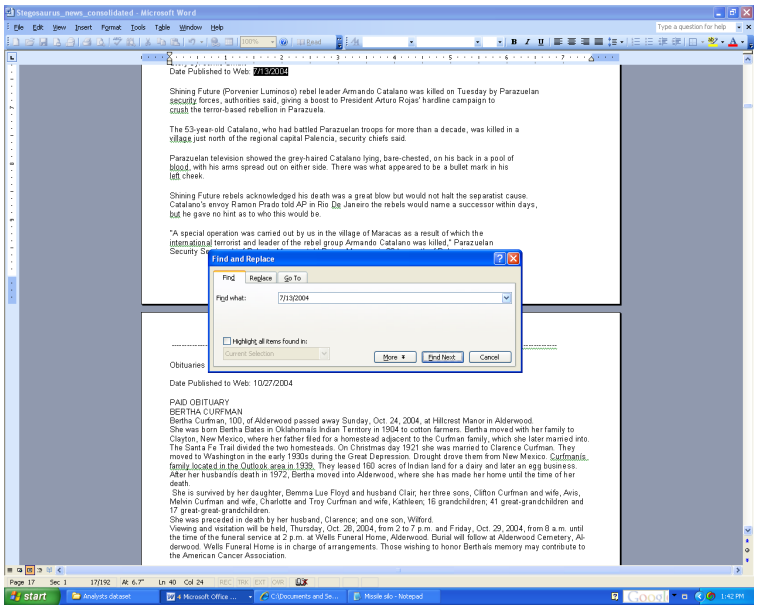

Figure B.5: Final state of the workspace for subject S2

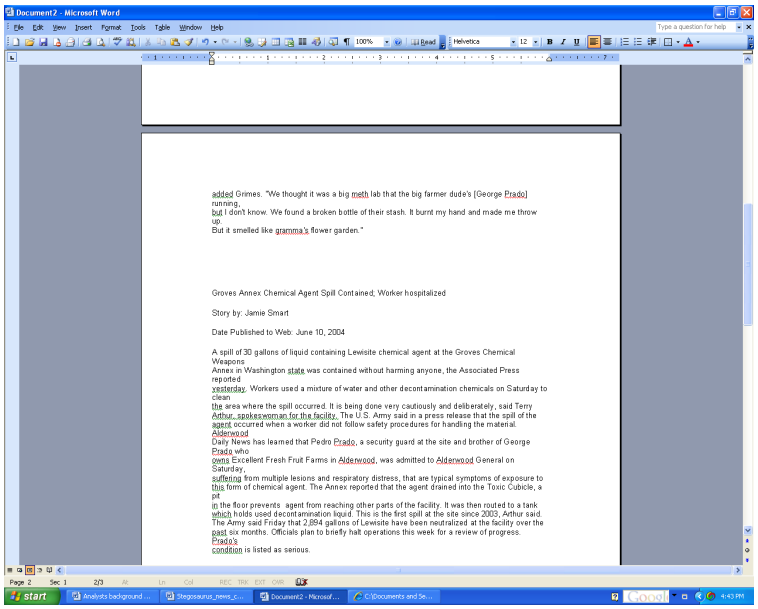

Figure B.6: Final state of the workspace for subject S3

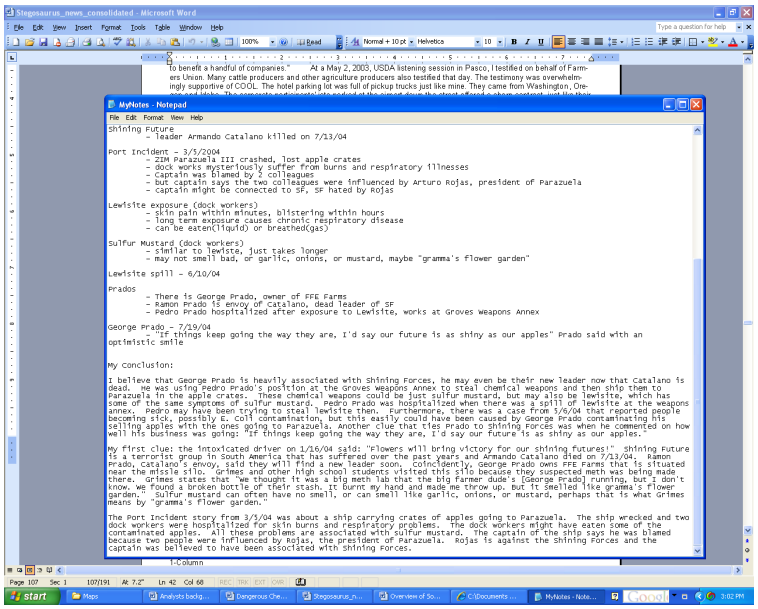

Figure B.7: Final state of the workspace for subject S4

Appendix C

Analyst Study Screenshots

In this appendix we list all of the final screenshots from the analyst study.

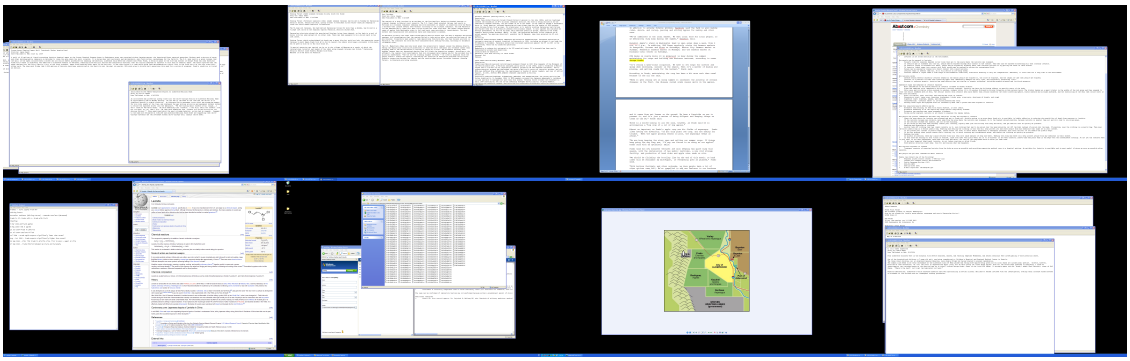

Figure C.1: Final state of the workspace for subject A1

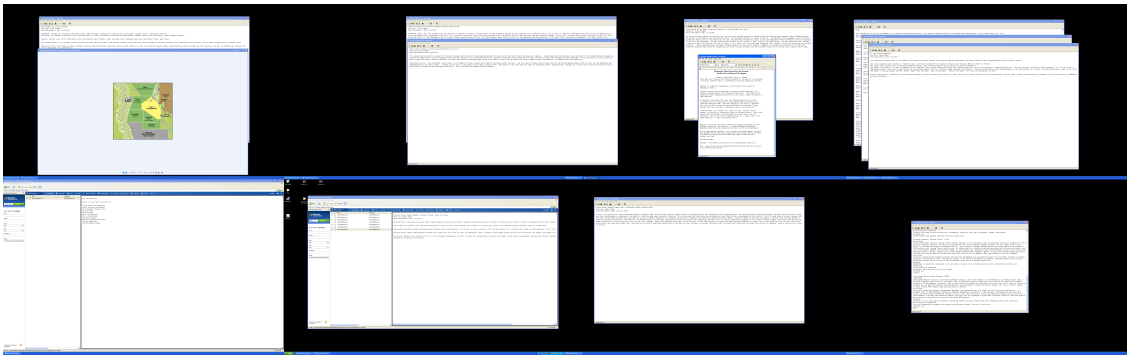

Figure C.2: Final state of the workspace for subject A2

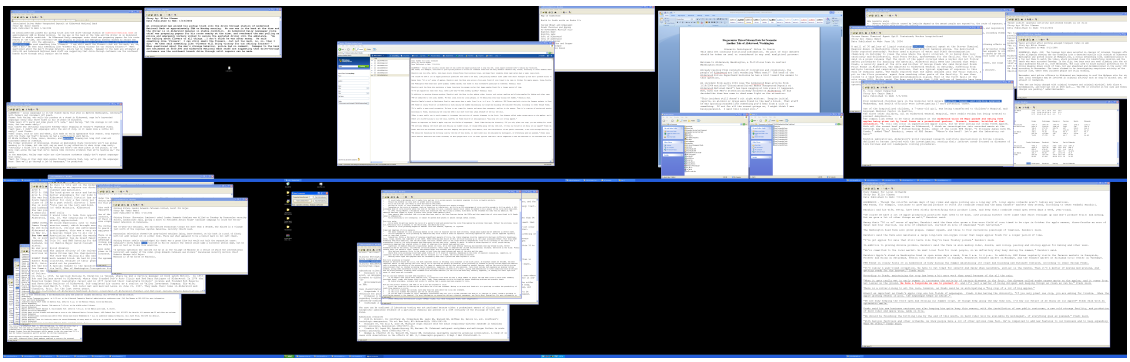

Figure C.3: Final state of the workspace for subject A3

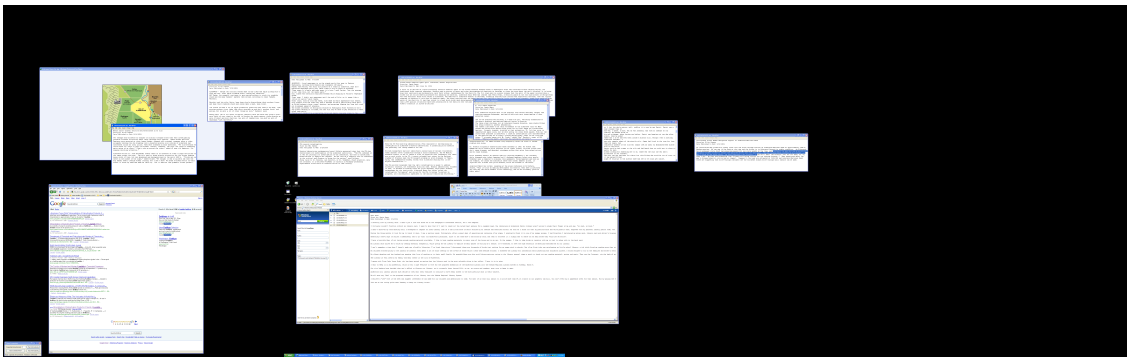

Figure C.4: Final state of the workspace for subject A4

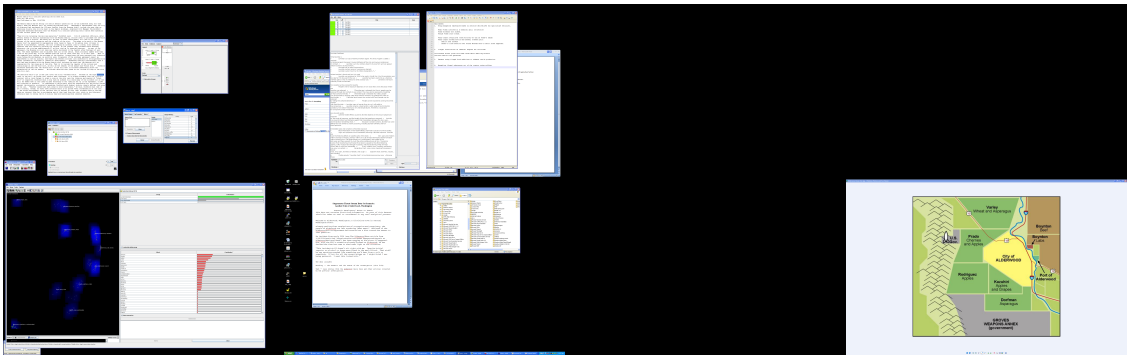

Figure C.5: Final state of the workspace for subject A5

Appendix D

Physical Navigation Study Documents

VirginiaTech

Office of Research Compliance
Institutional Review Board
2000 Kraft Drive, Suite 2000 (0497)
Blacksburg, Virginia 24060
540/231-4606 Fax 540/231-0959
e-mail irb@vt.edu
Website: www.irb.vt.edu

MEMORANDUM

DATE: July 28, 2010

TO: Christopher L. North, Christopher Andrews

FROM: Virginia Tech Institutional Review Board (FWA00000572, expires June 13, 2011)

PROTOCOL TITLE: Evaluating Spatial Persistence

IRB NUMBER: 10-622

Effective July 28, 2010, the Virginia Tech IRB Chair, Dr. David M. Moore, approved the new protocol for the above-mentioned research protocol.

This approval provides permission to begin the human subject activities outlined in the IRB-approved protocol and supporting documents.

Plans to deviate from the approved protocol and/or supporting documents must be submitted to the IRB as an amendment request and approved by the IRB prior to the implementation of any changes, regardless of how minor, except where necessary to eliminate apparent immediate hazards to the subjects. Report promptly to the IRB any injuries or other unanticipated or adverse events involving risks or harms to human research subjects or others.

All investigators (listed above) are required to comply with the researcher requirements outlined at http://www.irb.vt.edu/pages/responsibilities.htm (please review before the commencement of your research).

PROTOCOL INFORMATION:
Approved as: **Expedited, under 45 CFR 46.110 category(ies) 6, 7**
Protocol Approval Date: **7/28/2010**
Protocol Expiration Date: **7/27/2011**
Continuing Review Due Date*: **7/13/2011**
*Date a Continuing Review application is due to the IRB office if human subject activities covered under this protocol, including data analysis, are to continue beyond the Protocol Expiration Date.

FEDERALLY FUNDED RESEARCH REQUIREMENTS:
Per federally regulations, 45 CFR 46.103(f), the IRB is required to compare all federally funded grant proposals / work statements to the IRB protocol(s) which cover the human research activities included in the proposal / work statement before funds are released. Note that this requirement does not apply to Exempt and Interim IRB protocols, or grants for which VT is not the primary awardee.

The table on the following page indicates whether grant proposals are related to this IRB protocol, and which of the listed proposals, if any, have been compared to this IRB protocol, if required.

Informed Consent for Participant of Investigative Project
Virginia Polytechnic Institute and State University

Title of Project: Evaluating Spatial Persistence

Principal Investigators: Christopher Andrews, Chris North

I. THE PURPOSE OF THIS RESEARCH/PROJECT

The purpose of this project is to examine the role of spatial persistence in sensemaking tasks. The results of this study will inform the development of future analytic tools.

II. PROCEDURES

You will be asked to analyze a collection of documents using a simple space-based analytic tool. The task will involve reading text documents, organizing them into the spatial environment, and analyzing them. At the end of the session, you will be asked to provide a report on your interpretation of the important information contained in the dataset.

The entire session will be video tapped, your activity with the tool captured, and you will occasionally be interrupted and asked about your progress.

III. RISKS

There is minimal risk involved with this study. The proposed experiment involves performing everyday work related tasks on a large, high-resolution display. The only foreseeable risk is mild mental fatigue caused by performing the tasks. If you experience headaches or other discomforts, please feel free to step back from the display and inform the experimenter.

IV. BENEFITS OF THIS PROJECT

Your participation in this project will provide information that may be used to inform the design of tools and applications for large displays. No guarantee of benefits has been made to encourage you to participate. You may receive a synopsis summarizing this research when completed. If you are interested, leave an email address with the investigator.

You are requested to refrain from discussing the evaluation with other people who might be in the candidate pool from which other participants might be drawn.

V. COMPENSATION

You will be compensated $5/hour for participating in this study, or $15 for completing the session. If you decide to leave early, you will be paid for the time you were participating. In addition, you have the opportunity to win an additional $25 by providing the highest scoring solution in your condition.

VI. EXTENT OF ANONYMITY AND CONFIDENTIALITY

The results of this study will be kept strictly confidential. Your written consent is required for the researchers to release any data identified with you as an individual to anyone other than personnel working on the project. While a video and audio record will be made, only unique subject numbers will be used in analysis and reports.

VIII. FREEDOM TO WITHDRAW

You are free to withdraw from this study at any time for any reason.

IX. APPROVAL OF RESEARCH

This research has been approved, as required, by the Institutional Review Board for projects involving human subjects at Virginia Polytechnic Institute and State University, and by the Department of Computer Science.

X. SUBJECT'S RESPONSIBILITIES AND PERMISSION

I voluntarily agree to participate in this study, and I know of no reason I cannot participate. I have read and understand the informed consent and conditions of this project. I have had all my questions answered. I hereby acknowledge the above and give my voluntary consent for participation in this project. If I participate, I may withdraw at any time without penalty. I agree to abide by the rules of this project

_____ _____

Signature Date

_____ _____

Name (please print) Contact: phone or address or

_____

email address (OPTIONAL)

Should I have any questions about this research or its conduct, I may contact:

Investigator: Christopher Andrews Phone 978.314.7881
email: cpa@cs.vt.edu

Review Board: David M. Moore, Office of Research Compliance, CVM Phase II (0442) 231-4991

cc: the participant, Christopher Andrews

Subject ID: _____

Demographic Information

Gender :

Age:

Occupation (if student, indicate graduate or undergraduate):

Major/ Area of Specialization

Describe any experience with intelligence analysis-type problems:

Describe any experience with large displays or multiple monitor systems:

1. WHO: who are the players engaging in questionable activities in the plot(s)? When appropriate, specify the association/organization they are associated with.

Name	Associated organization	Involved in illegal activities? (Yes/No)	Involved in terrorist activities? (Yes/No)	Most relevant source files (5 MAX)	Notes
e.g. John doe	Virginia Tech	Yes	Yes	FBI_1, FBI_2	

(Note: ignore the file extensions of the files)

2. WHEN /WHAT: What events occurred during this time frame that are most relevant to the plot(s)?

Provide a text list of events following the sample layout. Use short description (i.e. one or 2 lines per event).
Provide what you think is the best subset of events (20 events MAX)

	Date (Can be a range)	**Event description**	**Most relevance source files (5 Max)**
	e.g. 11/4/1789	e.g. Citizens take control of the Bastille	CIA_4, Image 2
1			
2			
3			
4			
5			
6			
7			
8			
9			
10			
11			
12			
13			
14			
15			
16			
17			
18			
19			
20 max			

3. WHERE: What locations are most relevant to the plot(s)?

Follow this example layout. Use only one-line per item.

	Location	**Description**	**Most relevance source files (5 Max)**
	e.g. Berlin	e.g. training of the bank robbers	
1			
2			
3			
4			
5			
6			
7			
8			

4. DEBRIEFING

Analysts will typically have to produce a report at the end of their investigation. This report is essentially a summary of the above. Include your written assessment of the situation (between **1000 and 2000 words**)

This narrative should describe

- the plot(s) and subplots(s) and how people, motivations, activities and locations are part of the plot.
- the relationships of the various players.
- If there are uncertainties, you can suggest possible next steps to clarify those uncertainties.
- Be sure to indicate which parts of your report you are certain versus uncertain about and why.

(NOTE: Focus on demonstrating why you believe what you believe)

Appendix E

Physical Navigtion Study Screenshots

In this appendix we list all of the final screenshots from the navigation comparison study.

Large display users

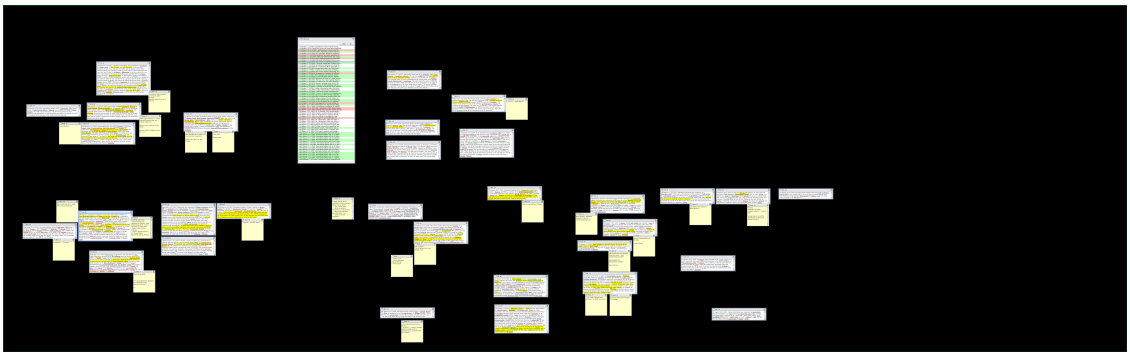

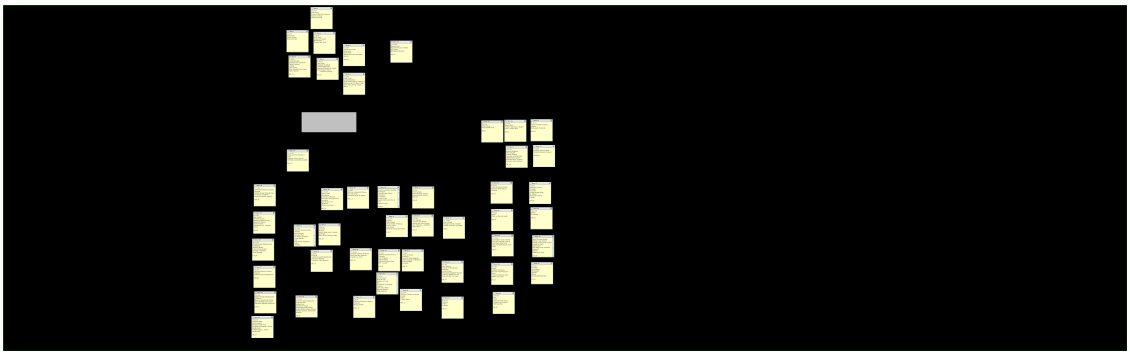

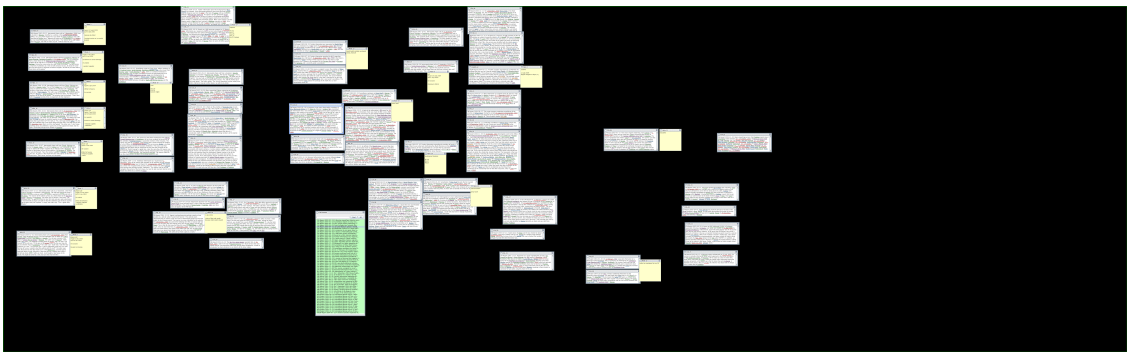

Small display users

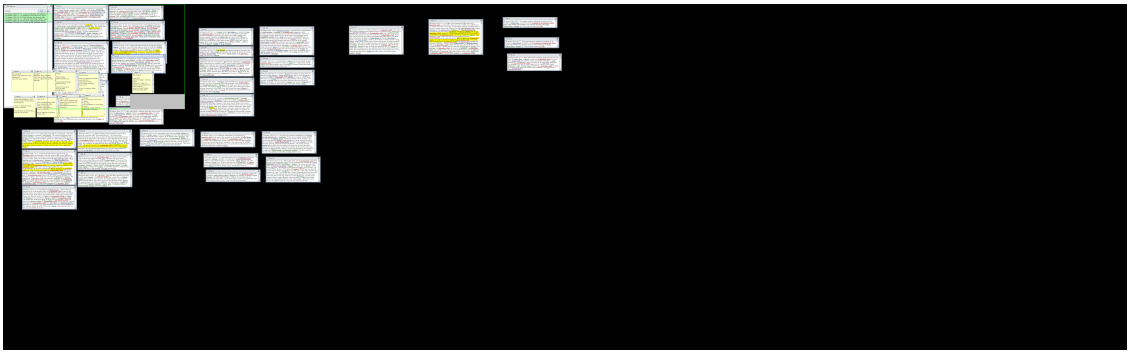

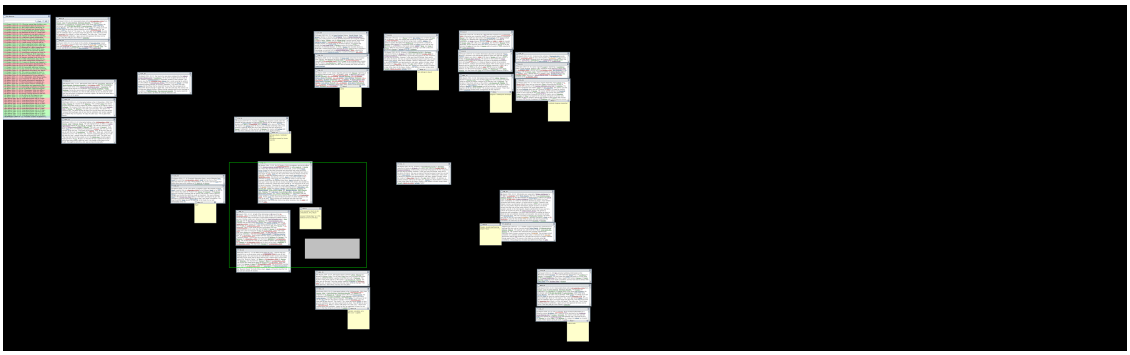

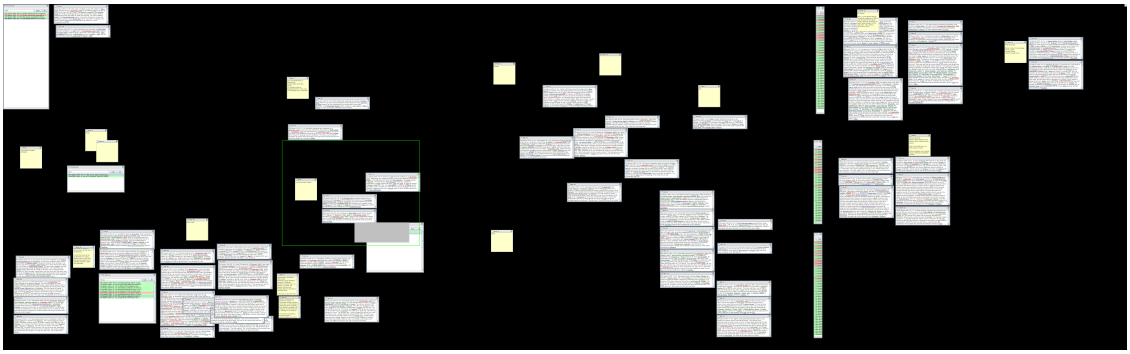

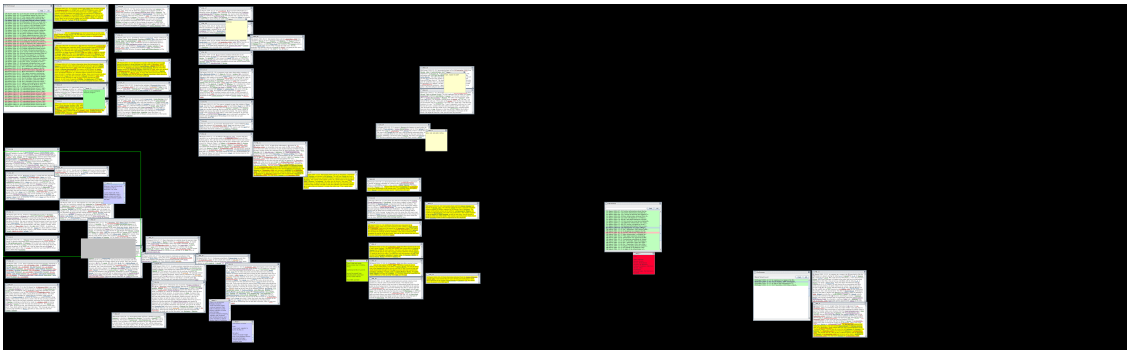

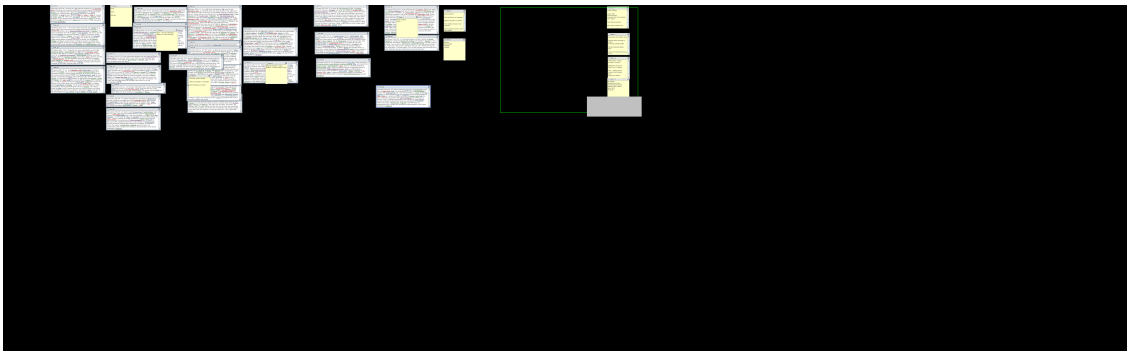

Bibliography

Analyst's Notebook, fetched July 2011. URL http://www.i2group.com.

Palantir, fetched July 2011. URL http://www.palantirtech.com/.

AlchemyAPI, fetched Aug 2011. URL http://www.alchemyapi.com/.

Lingpipe, fetched July 2011. URL http://alias-i.com/lingpipe/.

Apache Lucene, fetched July 2011. URL http://lucene.apache.org.

Sentinel Visualizer, fetched July 2011. URL http://www.fmsasg.com.

C. Andrews, A. Endert, and C. North. Space to think: large high-resolution displays for sensemaking. In *Proceedings of the 28th international conference on Human factors in computing systems*, CHI '10, pages 55–64, New York, NY, USA, 2010. ACM.

C. Andrews, M. S. Hossain, S. Gad, N. Ramakrishnan, and C. North. Analyst's Workspace: Protecting Vastopolis - VAST 2011 Mini Challenge 3 award: "Novel use of large screen workspace to support analysis". In *VAST '11*, to appear.

M. N. Avraamides and S. G. Sofroniou. Spatial frameworks in imagined navigation. *Psychon Bull Rev*, 13(3):510–515, 2006.

R. Ball and C. North. The effects of peripheral vision and physical navigation on large scale visualization. In *GI '08: Proceedings of graphics interface 2008*, pages 9–16, Toronto, Ont., Canada, Canada, 2008. Canadian Information Processing Society.

R. Ball, C. North, and D. A. Bowman. Move to improve: promoting physical navigation to increase user performance with large displays. In *CHI '07: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 191–200, New York, NY, USA, 2007. ACM Press.

P. Baudisch, N. Good, V. Bellotti, and P. Schraedley. Keeping things in context: a comparative evaluation of focus plus context screens, overviews, and zooming. In *CHI '02: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 259–266, New York, NY, USA, 2002. ACM Press.

D. Bauer, P. Fastrez, and J. Hollan. Spatial tools for managing personal information collections. *System Sciences, 2005. HICSS '05. Proceedings of the 38th Annual Hawaii International Conference on*, pages 104b–104b, Jan. 2005.

B. B. Bederson and A. Boltman. Does animation help users build mental maps of spatial information? In *INFOVIS '99: Proceedings of the 1999 IEEE Symposium on Information Visualization*, page 28, Washington, DC, USA, 1999. IEEE Computer Society.

B. B. Bederson, L. Stead, and J. D. Hollan. Pad++: advances in multiscale interfaces. In *CHI '94: Conference companion on Human factors in computing systems*, pages 315–316, New York, NY, USA, 1994. ACM.

M. Bernstein. Collage, composites, construction. In *HYPERTEXT '03: Proceedings of the fourteenth ACM conference on Hypertext and hypermedia*, pages 122–123, New York, NY, USA, 2003. ACM.

X. Bi and R. Balakrishnan. Comparing usage of a large high-resolution display to single or dual desktop displays for daily work. In *CHI '09: Proceedings of the 27th international conference on Human factors in computing systems*, pages 1005–1014, New York, NY, USA, 2009. ACM.

E. A. Bier, E. W. Ishak, and E. Chi. Entity workspace: An evidence file that aids memory, inference, and reading. In *Intelligence and Security Informatics 2006*, pages 466–472, 2006.

E. A. Bier, S. K. Card, and J. W. Bodnar. Entity-based collaboration tools for intelligence analysis. *Visual Analytics Science and Technology, 2008. VAST '08. IEEE Symposium on*, pages 99–106, Oct. 2008.

F. Biocca, A. Tang, D. Lamas, J. Gregg, P. Gai, and R. Brady. How do users organize virtual tools around their body in immersive virtual and augmented reality environments. *An exploratory study of egocentric spatial mappings of virtual tools in the mobile infosphere. East Lansing, MI: Media Interface and Network Design Lab*, 2001.

T. Blanc-Brude and D. L. Scapin. What do people recall about their documents?: implications for desktop search tools. In *IUI '07: Proceedings of the 12th international conference on Intelligent user interfaces*, pages 102–111, New York, NY, USA, 2007. ACM.

T. Boren and J. Ramey. Thinking aloud: reconciling theory and practice. *Professional Communication, IEEE Transactions on*, 43(3):261–278, January 2000.

A. Clark. An embodied cognitive science. *Trends in Cognitive Sciences*, 3(9):345–351, Sept 1999.

A. Clark. *Supersizing the Mind: Embodiment, Action, and Cognitive Extension.* Oxford University Press, New York, 2008.

A. Cockburn. Revisiting 2d vs 3d implications on spatial memory. In *AUIC '04: Proceedings of the fifth conference on Australasian user interface*, pages 25–31, Darlinghurst, Australia, Australia, 2004. Australian Computer Society, Inc.

A. Cockburn and B. McKenzie. Evaluating the effectiveness of spatial memory in 2d and 3d physical and virtual environments. In *CHI '02: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 203–210, New York, NY, USA, 2002. ACM.

A. Cockburn, A. Karlson, and B. B. Bederson. A review of overview+detail, zooming, and focus+context interfaces. *ACM Comput. Surv.*, 41:2:1–2:31, January 2009.

P. Consulting. The 30-inch apple cinema hd display productivity benchmark. Technical report, Pfeiffer Consulting, 2005.

L. Costello, G. Grinstein, C. Plaisant, and J. Scholtz. Advancing user-centered evaluation of visual analytic environments through contests. *Information Visualization*, 8:230–238, June 2009.

M. Czerwinski, D. S. Tan, and G. G. Robertson. Women take a wider view. In *CHI '02: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 195–202, New York, NY, USA, 2002. ACM.

M. P. Czerwinski, G. Smith, T. Regan, B. Meyers, G. G. Robertson, and G. Starkweather. Toward characterizing the productivity benefits of very large displays. In *Interact*, pages 9–16. IFIP, 2003.

J. D. Austin Henderson and S. Card. Rooms: the use of multiple virtual workspaces to reduce space contention in a window-based graphical user interface. *ACM Trans. Graph.*, 5(3):211–243, 1986.

R. P. Darken and J. L. Sibert. Wayfinding strategies and behaviors in large virtual worlds. In *CHI '96: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 142–149, New York, NY, USA, 1996. ACM Press.

P. Dourish. *Embodied Interaction: Exploring the Foundations of a New Approach to HCI (1999, draft overview paper, Ch5)*. MIT Press, 2001.

R. Eccles, T. Kapler, R. Harper, and W. Wright. Stories in geotime. *Visual Analytics Science and Technology, 2007. VAST 2007. IEEE Symposium on*, pages 19–26, 30 2007-Nov. 1 2007.

B. D. Ehret. Learning where to look: location learning in graphical user interfaces. In *CHI '02: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 211–218, New York, NY, USA, 2002. ACM.

A. Fass, J. Forlizzi, and R. Pausch. Messydesk and messyboard: two designs inspired by the goal of improving human memory. In *DIS '02: Proceedings of the 4th conference on Designing interactive systems*, pages 303–311, New York, NY, USA, 2002. ACM.

L. Francisco-Revilla and F. Shipman. Parsing and interpreting ambiguous structures in spatial hypermedia. In *HYPERTEXT '05: Proceedings of the sixteenth ACM conference on Hypertext and hypermedia*, pages 107–116, New York, NY, USA, 2005. ACM.

G. W. Furnas. Generalized fisheye views. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, CHI '86, pages 16–23, New York, NY, USA, 1986. ACM.

G. W. Furnas and B. B. Bederson. Space-scale diagrams: understanding multiscale interfaces. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, CHI '95, pages 234–241, New York, NY, USA, 1995. ACM Press/Addison-Wesley Publishing Co.

W. D. Gray and W.-T. Fu. Soft constraints in interactive behavior: the case of ignoring perfect knowledge in-the-world for imperfect knowledge in-the-head. *Cognitive Science: A Multidisciplinary Journal*, 28(3):359–382, 2004.

G. Grinstein, T. O'Connell, S. Laskowski, C. Plaisant, J. Scholtz, and M. Whiting. Vast 2006 contest ? a tale of alderwood. *Symposium On Visual Analytics Science And Technology*, 0:215–216, 2006.

J. Grudin. Partitioning digital worlds: focal and peripheral awareness in multiple monitor use. In *CHI '01: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 458–465, New York, NY, USA, 2001. ACM Press.

F. G. Halasz. Reflections on notecards: seven issues for the next generation of hypermedia systems. *Commun. ACM*, 31(7):836–852, 1988.

F. G. Halasz, T. P. Moran, and R. H. Trigg. Notecards in a nutshell. *SIGCHI Bull.*, 17(SI): 45–52, 1987.

S. M. Hess, M. C. Detweiler, and R. D. Ellis. The utility of display space in keeping track of rapidly changing information. *Human Factors*, 41(2):257–281, 1999.

R. J. Heuer. *Psychology of Intelligence Analysis*. Center for the Study of Intelligence, 1999.

J. Hollan, E. Hutchins, and D. Kirsh. Distributed cognition: toward a new foundation for human-computer interaction research. *ACM Trans. Comput.-Hum. Interact.*, 7(2):174–196, 2000.

M. S. Hossain, C. Andrews, N. Ramakrishnan, and C. North. Helping intelligence analysts make connections. In *AAAI '11*, to appear.

D. R. Hutchings. *Making multiple monitors more manageable*. PhD thesis, Georgia Tech, Atlanta, GA, USA, 2006. Adviser-John Stasko.

D. R. Hutchings, G. Smith, B. Meyers, M. Czerwinski, and G. Robertson. Display space usage and window management operation comparisons between single monitor and multiple monitor users. In *AVI '04: Proceedings of the working conference on Advanced visual interfaces*, pages 32–39, New York, NY, USA, 2004. ACM Press.

E. Hutchins. Distributed cognition. In N. J. Smelser and P. B. Baltes, editors, *International Encyclopedia of the Social & Behavioral Sciences*, pages 2068 – 2072. Pergamon, Oxford, 2001.

R. Johnston. Developing a taxonomy of intelligence analysis variables. *CIA Studies in Intelligence*, 47(3), 2003.

W. P. Jones and S. T. Dumais. The spatial metaphor for user interfaces: experimental tests of reference by location versus name. *ACM Trans. Inf. Syst.*, 4(1):42–63, 1986.

D. Jonker, W. Wright, D. Schroh, P. Proulx, and B. Cort. Information triage with trist. Technical report, Oculus, 2005.

Y. Kang, C. Gorg, and J. Stasko. Evaluating visual analytics systems for investigative analysis: Deriving design principles from a case study. In *Visual Analytics Science and Technology, 2009. VAST 2009. IEEE Symposium on*, pages 139 –146, 12-13 2009.

Y.-a. Kang and J. Stasko. Lightweight task/application performance using single versus multiple monitors: a comparative study. In *GI '08: Proceedings of graphics interface 2008*, pages 17–24, Toronto, Ont., Canada, Canada, 2008. Canadian Information Processing Society.

T. Kapler and W. Wright. Geotime information visualization. *Information Visualization*, 4 (2):136–146, 2005.

D. Kirsh. The intelligent use of space. *Artif. Intell.*, 73(1-2):31–68, 1995.

D. Kirsh and P. P. Maglio. On distinguishing epistemic from pragmatic action. *Cognitive Science*, 18(4):513–549, 1994.

W. Kuhn. Handling data spatially: Spatializating user interfaces. In M.-J. Kraak and M. Molenaar, editors, *Advances in GIS Research - 7th International Symposium on Spatial Data Handling (SDH'96)*, pages 877–893, 1996.

G. Lakoff and M. Johnson. *Metaphors We Live By.* University of Chicago Press, 1980.

Z. Liu, N. J. Nersessian, and J. T. Stasko. Distributed cognition as a theoretical framework for information visualization. In *IEEE Transactions on Visualization and Computer Graphics*, volume 14, 2008.

T. W. Malone. How do people organize their desks?: Implications for the design of office information systems. *ACM Trans. Inf. Syst.*, 1(1):99–112, 1983.

C. C. Marshall and I. Frank M. Shipman. Spatial hypertext and the practice of information triage. In *HYPERTEXT '97: Proceedings of the eighth ACM conference on Hypertext*, pages 124–133, New York, NY, USA, 1997. ACM.

C. C. Marshall and R. A. Rogers. Two years before the mist: experiences with aquanet. In *ECHT '92: Proceedings of the ACM conference on Hypertext*, pages 53–62, New York, NY, USA, 1992. ACM.

C. C. Marshall, F. G. Halasz, R. A. Rogers, and J. William C. Janssen. Aquanet: a hypertext tool to hold your knowledge in place. In *HYPERTEXT '91: Proceedings of the third annual ACM conference on Hypertext*, pages 261–275, New York, NY, USA, 1991. ACM.

K. Nakakoji, Y. Yamamoto, S. Takada, and B. N. Reeves. Two-dimensional spatial positioning as a means for reflection in design. In *DIS '00: Proceedings of the 3rd conference on Designing interactive systems*, pages 145–154, New York, NY, USA, 2000. ACM.

T. Ni, D. A. Bowman, and J. Chen. Increased display size and resolution improve task performance in information-rich virtual environments. In *GI '06: Proceedings of Graphics Interface 2006*, pages 139–146, Toronto, Ont., Canada, Canada, 2006. Canadian Information Processing Society.

D. A. Norman. Cognitive artifacts. *Cambridge Series On Human-Computer Interaction*, pages 17–38, 1991.

J. K. O'Regan. Solving the "real" mysteries of visual perception: the world as an outside memory. *Can J Psychol*, 46(3):461–488, Sep 1992.

J. Patten and H. Ishii. A comparison of spatial organization strategies in graphical and tangible user interfaces. In *DARE '00: Proceedings of DARE 2000 on Designing augmented reality environments*, pages 41–50, New York, NY, USA, 2000. ACM.

W. Pike, R. May, and A. Turner. Supporting knowledge transfer through decomposable reasoning artifacts. *System Sciences, 2007. HICSS 2007. 40th Annual Hawaii International Conference on*, pages 204c–204c, Jan. 2007.

P. Pirolli and S. Card. The sensemaking process and leverage points for analyst technology as identified through cognitive task analysis the sensemaking process and leverage points for analyst technology as identified through cognitive task analysis. In *International Conference on Intelligence Analysis*, McLean, VA, May 2005.

C. Plaisant, G. Grinstein, J. Scholtz, M. Whiting, T. O'Connell, S. Laskowski, L. Chien, A. Tat, W. Wright, C. Gorg, Z. Liu, N. Parekh, K. Singhal, and J. Stasko. Evaluating visual analytics at the 2007 vast symposium contest. *Computer Graphics and Applications, IEEE*, 28(2):12–21, March-April 2008.

M. D. Plumlee and C. Ware. Zooming versus multiple window interfaces: Cognitive costs of visual comparisons. *ACM Trans. Comput.-Hum. Interact.*, 13(2):179–209, 2006.

E. D. Ragan, A. Sowndararajan, R. Kopper, and D. A. Bowman. The effects of higher levels of immersion on procedure memorization performance and implications for educational virtual environments. *Presence: Teleoper. Virtual Environ.*, 19:527–543, 2010.

D. Richardson and M. Spivey. Representation, space and hollywood squares: Looking at things that aren't there anymore, September 2000.

G. Robertson, M. Czerwinski, K. Larson, D. C. Robbins, D. Thiel, and M. van Dantzich. Data mountain: using spatial memory for document management. In *UIST '98: Proceedings of the 11th annual ACM symposium on User interface software and technology*, pages 153–162, New York, NY, USA, 1998. ACM.

G. Robertson, M. Czerwinski, P. Baudisch, B. Meyers, D. Robbins, G. Smith, and D. Tan. The large-display user experience. *IEEE Comput. Graph. Appl.*, 25(4):44–51, 2005.

A. C. Robinson. *Design for Synthesis in Geovisualization*. PhD thesis, Pennsylvania State University, University Park, PA, August 2008a.

A. C. Robinson. Collaborative synthesis of visual analytic results. *Visual Analytics Science and Technology, 2008. VAST '08. IEEE Symposium on*, pages 67–74, Oct. 2008b.

M. Rønne Jakobsen and K. Hornbæk. Sizing up visualizations: effects of display size in focus+context, overview+detail, and zooming interfaces. In *Proceedings of the 2011 annual conference on Human factors in computing systems*, CHI '11, pages 1451–1460, New York, NY, USA, 2011. ACM.

D. Russell, M. Slaney, Y. Qu, and M. Houston. Being literate with large document collections: Observational studies and cost structure tradeoffs. In *System Sciences, 2006. HICSS '06.*

Proceedings of the 39th Annual Hawaii International Conference on, volume 3, pages 55–55, Jan. 2006.

D. M. Russell, M. J. Stefik, P. Pirolli, and S. K. Card. The cost structure of sensemaking. In *CHI '93: Proceedings of the INTERACT '93 and CHI '93 conference on Human factors in computing systems*, pages 269–276, New York, NY, USA, 1993. ACM.

A. J. Sabri, R. G. Ball, A. Fabian, S. Bhatia, and C. North. High-resolution gaming: Interfaces, notifications, and the user experience. *Interact. Comput.*, 19(2):151–166, 2007.

F. M. Shipman, III and C. C. Marshall. Spatial hypertext: an alternative to navigational and semantic links. *ACM Comput. Surv.*, page 14, 1999a.

F. M. Shipman, III and C. C. Marshall. Formality considered harmful: Experiences, emergingthemes, and directions on the use of formal representations ininteractive systems. *Comput. Supported Coop. Work*, 8(4):333–352, 1999b.

F. M. Shipman, III, C. C. Marshall, and T. P. Moran. Finding and using implicit structure in human-organized spatial layouts of information. In *CHI '95: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 346–353, New York, NY, USA, 1995. ACM Press/Addison-Wesley Publishing Co.

F. M. Shipman, III, H. Hsieh, R. Airhart, P. Maloor, J. M. Moore, and D. Shah. Emergent structure in analytic workspaces: Design and use of the visual knowledge builder. In *Proceedings of IFIP INTERACT01: Human-Computer Interaction*, pages 132–139. IFIP Technical Committee No 13 on Human-Computer Interaction, 2001a.

F. M. Shipman, III, H. Hsieh, P. Maloor, and J. M. Moore. The visual knowledge builder: a second generation spatial hypertext. In *HYPERTEXT '01: Proceedings of the 12th ACM*

conference on Hypertext and Hypermedia, pages 113–122, New York, NY, USA, 2001b. ACM.

F. M. Shipman, III, J. M. Moore, P. Maloor, H. Hsieh, and R. Akkapeddi. Semantics happen: knowledge building in spatial hypertext. In *HYPERTEXT '02: Proceedings of the thirteenth ACM conference on Hypertext and hypermedia*, pages 25–34, New York, NY, USA, 2002. ACM.

L. Shupp, C. Andrews, M. Dickey-Kurdziolek, B. Yost, and C. North. Shaping the display of the future: The effects of display size and curvature on user performance and insights. *Human-Computer Interaction*, 24(1&2):230 – 272, January 2009.

J. Stasko, C. Gorg, Z. Liu, and K. Singhal. Jigsaw: Supporting investigative analysis through interactive visualization. *Visual Analytics Science and Technology, 2007. VAST 2007. IEEE Symposium on*, pages 131–138, 30 2007-Nov. 1 2007.

K. Swaminathan and S. Sato. Interaction design for large displays. *interactions*, 4(1):15–24, 1997.

D. Tan. *Exploiting the Cognitive and Social Benefits of Physically Large Displays*. PhD thesis, Carnegie Mellon University, 2004.

D. S. Tan, D. Gergle, P. Scupelli, and R. Pausch. Physically large displays improve performance on spatial tasks. *ACM Trans. Comput.-Hum. Interact.*, 13(1):71–99, 2006.

J. J. Thomas and K. A. Cook, editors. *Illuminating the Path: The Research and Development Agenda for Visual Analytics*. National Visualization and Analytics Center, 2005.

M. Waldner, W. Puff, A. Lex, M. Streit, and D. Schmalstieg. Visual links across applications. In *Proceedings of Graphics Interface 2010*, GI '10, pages 129–136, Toronto, Ont., Canada, Canada, 2010. Canadian Information Processing Society.

J. A. Wise, J. J. Thomas, K. Pennock, D. Lantrip, M. Pottier, A. Schur, and V. Crow. Visualizing the non-visual: spatial analysis and interaction with information for text documents. In *INFOVIS '95*, pages 51–58. IEEE Computer Society, 1995.

W. Wright, D. Schroh, P. Proulx, A. Skaburskis, and B. Cort. The sandbox for analysis: concepts and methods. In *CHI '06: Proceedings of the SIGCHI conference on Human Factors in computing systems*, pages 801–810, New York, NY, USA, 2006. ACM.

Y. Yamamoto, K. Nakakoji, Y. Nishinaka, M. Asada, and R. Matsuda. What is the space for?: the role of space in authoring hypertext representations. In *HYPERTEXT '05: Proceedings of the sixteenth ACM conference on Hypertext and hypermedia*, pages 117–125, New York, NY, USA, 2005. ACM.

B. Yost, Y. Haciahmetoglu, and C. North. Beyond visual acuity: the perceptual scalability of information visualizations for large displays. In *CHI '07: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 101–110, New York, NY, USA, 2007. ACM.

J. Zhang. The nature of external representations in problem solving. *Cognitive Science*, 21: 179–217, 1997.

J. Zhang and D. Norman. Representations in Distributed Cognitive Tasks. *Cognitive Science*, 18(1):87–122, 1994.