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

Christopher Andrews, Alex Endert, and Chris North
Department of Computer Science
Virginia Polytechnic Institute and State University
Blacksburg, VA 24060, USA
{cpa, aendert, north}@cs.vt.edu

ABSTRACT
Space supports human cognitive abilities in a myriad of ways. The note attached to the side of the monitor, the papers spread out on the desk, diagrams scrawled on a whiteboard, and even the keys left out on the counter are all examples of using space to recall, reveal relationships, and think. Technological advances have made it possible to construct large display environments in which space has real meaning. This paper examines how increased space affects the way displays are regarded and used within the context of the cognitively demanding task of sensemaking. A pair of studies were conducted demonstrating how the spatial environment supports sensemaking by becoming part of the distributed cognitive process, providing both external memory and a semantic layer.

Author Keywords
large, high-resolution displays, visual analytics, space, sensemaking, distributed cognition, embodiment.

ACM Classification Keywords
H5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

General Terms
Experimentation

INTRODUCTION
In the last decade, we have seen large changes in display technology. It is now possible to construct displays that are not just physically large, but high-resolution as well (Figure 1). The result is a transformation in which the workspace takes on the characteristics of real space. Multiple objects can be displayed with full resolution and can have real spatial relationship, not only amongst themselves, but with respect to the user as well. What does this change in the spatiality mean for the user beyond the simple increase in the amount of information that can be viewed?

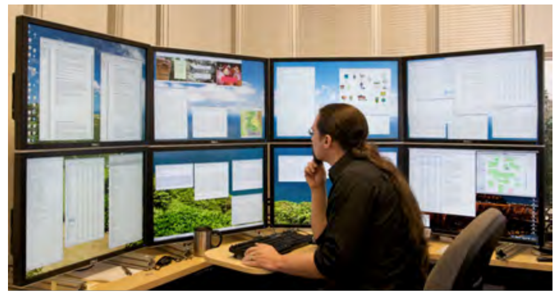

Figure 1 Prototype 10,240 x 3200 "analyst's workstation"

Our interest is in sensemaking, so our core question is: How can the space afforded by large, high-resolution displays support sensemaking? In other words, how can this space be used to think?

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. It is involved in problems that range from 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. Pirolli and Card have identified a number of "leverage points" in this process – problematic points in the process that would benefit from improvements [1]. These are primarily concerned with issues of data overload and attention management. Our explorations of the spatial characteristics of large, high-resolution displays directly address some of these problem areas, with particular emphasis on the use of the space for organization and memory.

LARGE, HIGH-RESOLUTION DISPLAYS
At their simplest, large, high-resolution displays allow users increased simultaneous access to information. The typical 17" monitor covers only about 10% of the visual field and only about 1% of what we can see by moving the head [2]. 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 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 this 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. A number of studies have illustrated these advantages, but have largely focused on relatively simple tasks (e.g., [3-5]).

It is our contention that large, high-resolution displays change more than just the amount of information that can be displayed – virtual and physical space merge to create a virtual workspace in which spatial relationships have real meaning. Other research has shown how the physical size of these displays alters how users regard them and how it affects the spatial perceptions of the user. Czerwinski et al. demonstrated that larger displays afford a greater field of view, significantly helping women to perform spatial tasks [6]. Tan completely isolated the size of the display, varying the distance from the user to maintain a constant resolution and viewing angle, and showed that large displays biased users towards adopting an egocentric perspective [7].

The increased number of available pixels also has the potential to create space by changing the dynamic between representation and position for objects displayed on them. 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. If we consider "typical" documents (e.g., text documents, web pages, images, presentations, etc.), we can represent them in a variety of ways. The "normal" or detailed view is one in which the details of the document can be directly perceived (i.e., text can be read). While this view depends on factors like font size and cropping it is generally quite space filling. Other representations include thumbnails, icons, and labels. Each of these requires progressively less room, but at the same time convey less information. The result of this dynamic is that on a conventional display, documents can either be arranged spatially or viewed at a detailed level – not both. On a large, high-resolution display, however, even the detailed view of a document is relatively small in relation to the available space, thus allowing documents to be placed in spatially meaningful ways while retaining their detailed representations.

In this environment, documents can be placed in a *persistent* location in space, allowing the user to navigate the document space using *physical navigation* [3] rather than having to switch between application level tasks (e.g., reading, annotating, search) and system level tasks (e.g., view management). The environment also allows the user to switch fluidly among multiple levels of representation. Details within the documents are available and dealt with as atomic units. The user can also interact with the document itself, considering it as a whole. In addition, the size space ratio makes it possible to arrange documents into patterns and structures that can be apprehended as a unit (e.g., a collection of related documents). In a large, high-resolution environment, all of these levels can be simultaneously available, creating a multi-scale document space.

Large, high-resolution displays are currently unique in their combination of physical affordances and display properties. Alternative display technologies such as tabletops, head-mounted displays (HMDs), CAVEs and even large projection screens provide environments where the physical properties of the display are meaningful, but they lack the ability to display the same density of data. Either the resolution is so low that document interaction is completely impractical (e.g., HMDs), or other techniques (e.g., pan and zoom) are required to navigate the space, reducing the physicality of the environment. Large, high-resolution displays, however, create an environment that blurs the lines between physical and virtual space, allowing greater opportunities for the perception and abilities we depend on in the physical world to be brought to bear.

THE STUDIES

In order to explore the role of space and large, high-resolution displays in the sensemaking process, we asked a number of users to solve an analytic problem using the facilities of our "analyst's workstation".

In many ways, these studies closely resemble Robinson's examination of analysts performing synthesis using physical artifacts [8]. 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 environment allows us to provide more technological support, thus broadening the focus to the whole sensemaking process.

The goals of these studies were twofold. The first goal of the study 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 [9] or Visual Knowledge Builder (VKB) [10].

The second goal was to examine the actual mechanics of sensemaking in this environment. In other words, we were interested in how analysts would use the space at the various stages of sensemaking, and how it affected the overall sensemaking process.

Study Details

The analytic problem presented to the participants was a data set originally developed for the interactive session of the VAST 2006 contest [11]. The dataset consists of 230 simulated news stories, three images (simulated advertisements), one map, one spreadsheet, and three reference documents. Of the 238 documents only about ten are required to establish ground truth. Several others provide some background information that is not necessary to solve the scenario, and the remainder is "noise".

There are several benefits to using this scenario. The solution is not trivial, but the problem was 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. There is also a known ground truth, which makes it easier to evaluate progress and the depth of insights arrived at by study participants.

In the first study, the data set was broken up into three different directories: news stories, background data, and images. In addition, these subjects were provided with a single document that contained all of the news stories consolidated into a single file. 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. For reasons discussed in the following sections, in the second study we removed the consolidated document and combined the simulated news articles and the background documents into a single directory.

The Analyst's Workstation

Our "analyst's workstation" is a large, 32 megapixel (10,240x3200) display. The display consists of a 4x2 grid of 30" LCD panels, each with a maximum resolution of 2560x1600 (Figure 1). Each column is freestanding, allowing the display to be curved around the user; this creates a natural center point at which the user is roughly equidistant from all of the tiles. A keyboard and mouse are placed on a rolling stand. A rolling chair was also provided, allowing subjects to move around the space.

The entire display is run from a single dedicated machine that was running Windows XP for this study. The benefit of the single-headed large display is that no special tools are required and standard applications work on it without modification. There are some limitations due to the limited support in window managers for large workspaces. For example, losing the cursor and windows and dialog boxes opening or gaining focus in unexpected locations are well known problems on larger displays [12], and will need to be addressed in the development of any future tools designed for spatial environments such as this one.

No special tools were developed for this study. With one exception that will be discussed later, all of the study participants only used basic text editors (e.g., WordPad), web browsers (Web use was not required, but some participants still consulted it), image viewers (e.g., Windows Picture and Fax Viewer), and Windows Explorer for file browsing and search.

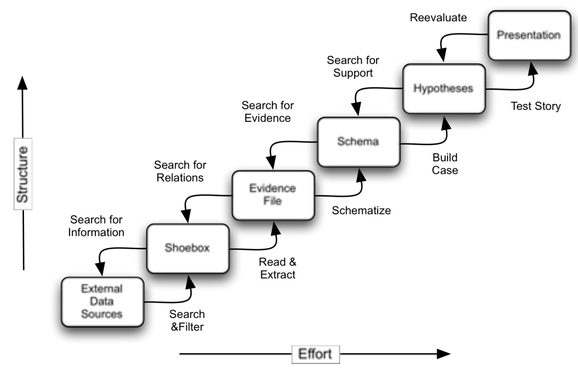

Figure 2 Sensemaking loop, adapted from Pirolli & Card [1]

Methodology

We approach these studies from the perspective of distributed cognition. We rely on two precepts from distributed cognition. First, humans are embodied, that is that human bodies and abilities and their interactions with the material world play a central role in cognition. The second, which follows from the first, is that cognition is distributed. In other words, we can consider individuals working with tools as comprising a *cognitive system* [13]. The focus of analysis based on distributed cognition is on the observable representations that are created and transformed by this cognitive system, helping us to understand their roles in the cognitive process.

Given this approach, the primary target of evaluation was the generation and transformation of representations in the environment. Where possible we drew connections back to Pirolli and Card's sensemaking loop [1] (Figure 2). The advantage of using this particular formulation of the sensemaking process as a framework is that it lays out the process as a series of stages, each with an associated structure that can be associated with the various representations generated by the analysts.

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 [14]. 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?

COMPARATIVE STUDY

The first study was a comparative study that compared the use the "analyst's workstation" with a conventional 17" monitor. Eight computer science graduate students were recruited to solve the scenario. Half of the subjects were given the use of the large display (L1-L4), while the rest

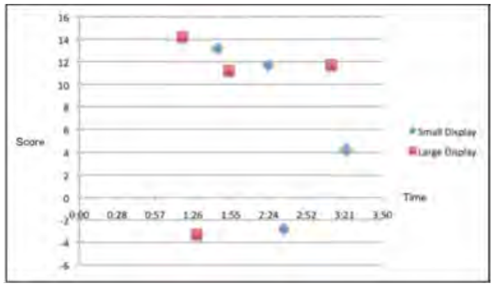

Figure 3 Quantitative results of the pilot study

used the conventional setup (S1-S4). Both groups were also provided the use of scrap paper and a whiteboard if they desired.

In addition to the analysis of the representations produced by these two groups, the sessions were timed and the results scored using a modified version of the metric proposed by Plaisant et al. [11]. The metric builds a composite score for a report based on the correct identification of participants, innocent bystanders, locations of interest, and motives. Subjects were given a maximum of four hours to work on the problem, with the fastest solution taking just under 1.5 hours and the longest solution taking close to 3.5 hours. Similarly, the presented solutions ranged in correctness from correct identification of ground truth with some fractions of a point subtracted for not clearly identifying true negatives (elements that are correctly reported as not being part of the plot) to negative scores resulting from identification of innocent bystanders and events as the center of the simulated plot. Given the small subject pool, the varied backgrounds, and the complexity of the problem, it is unsurprising that there were no significant differences between the two groups (Figure 3).

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.

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. 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.

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 [3].

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 [1]. 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.

This study revealed a number of interesting aspects of the use of space for sensemaking, but the use of the consolidated document limited this use primarily to supporting memory.

PROFESSIONAL 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 in 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. In addition, as described above, we reorganized the data set. We removed the consolidated news document to encourage an increased number of open documents, providing more opportunities for spatial strategies to develop. In the previous study, several subjects started by reading through the background documents before they began, which, because of their specific nature, biased the investigation towards chemical weapons and South American terrorist groups. To reduce this bias, we combined all of the documents together into a single data directory so that they would only come up through the result of a search.

Overall Reaction

Reaction to the display was overwhelmingly positive. Subject A1 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, 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 spatial environment had a clear impact on how the subjects approached the problem.

Non-spatial use of the display

The exception was A5, who, after performing one search with Windows Explorer, discovered a copy of IN-SPIRE [15] on the experimental machine and proceeded to use it for the rest of the session. 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 his limited use of the available space. In total, he used less than half of the display.

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. While IN-SPIRE's integrated query tool / document viewer only permits the user to view the contents of a single document at a time, the subject could have hunted down the original file and opened it separately, but that would have required extra work. 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 constrain the discussion to 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 our 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 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.

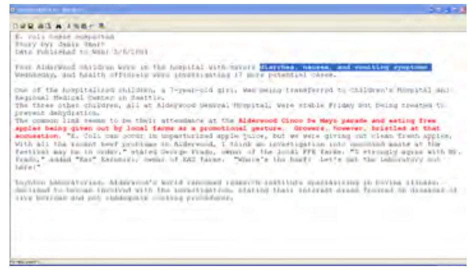

Figure 4 Example document showing both selection highlighting and text color highlighting

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 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 comparative 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. We

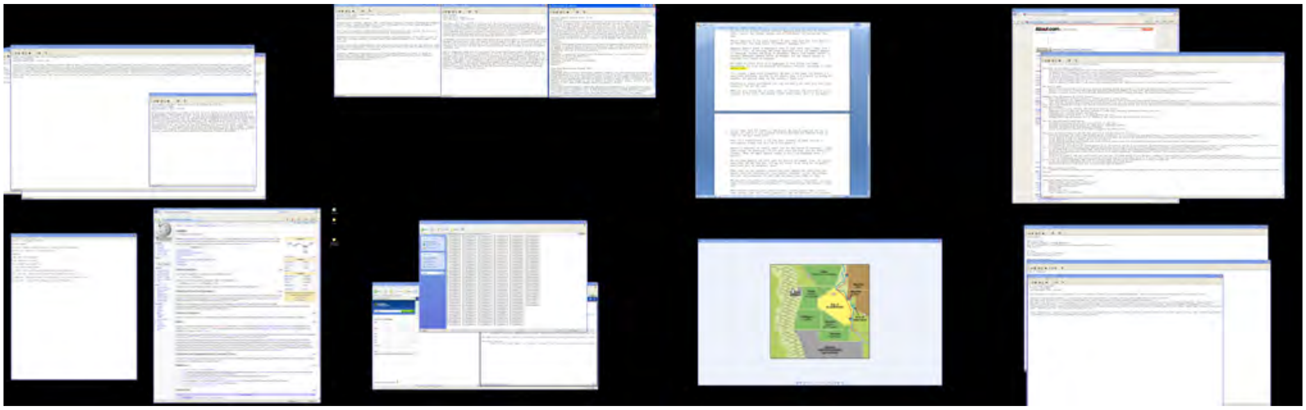

Figure 5 Final workspace for subject A1. Top row: 1) critical docs, 2) Shining Future, 3) George Prado, 4) chemical weapons. Bottom row: 5) timeline and wiki page, 6) work zone, 7) map, 8) potentials

are more interested in how the analysts made use of the surrounding area.

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 marshalling* and *schematizing*. The exceptions were small, localized organizations that arose as artifacts of the process – linear orderings that reflected the order in which documents were found. While not part of the marshalling 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 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. This behavior is particularly interesting because it is a common technique for making sense of physical, rather than digital, documents [16], showing how the environment has changed the analyst’s perception of the documents and the workspace.

An interesting characteristic of the clusters is that some were carefully arranged to maintain visibility, while others (like the “potentially interesting” cluster) were allowed to accumulate in a pile, with just enough left showing to trigger recognition. 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.

There are a couple of reasons why clustering formed the primary spatial metaphor. Categorization is one of the most basic forms of marshalling. 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 [17]. It is also low cost and relatively easy.

However, the structural properties of the display must also be acknowledged as playing a role in this direction. Monitor boundaries create natural partitions [2]. 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.

Regardless, clustering was not the only organizational structure that was observed. One interesting structure was a 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 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 6).

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

Figure 6 Subjects A4's timeline before (top) and after (bottom) restructuring to encode categorizations within the structure

throughout the process. We identified two distinct levels of integration: 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 marshalling. 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 marshalling 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 in the "analyst's workstation" provides the primary facilitation for the integration we observed.

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 "analyst's workstation", 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. 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., spatially oriented), was observed across multiple users, and showed adaptation to the environment. Examining these behaviors, we can classify the uses of space for sensemaking into two categories: memory and semantics.

Memory

Many of the representations humans create are ultimately a form of external memory. This works because objects in space are visible and *persistent*. It is important to consider not just physical or virtual artifacts, but the space they occupy as well. The to-do notes stuck to the monitor are as much about the location as they are about the contents. By being spatially persistent, and in a location of attention, the note facilitates opportunistic refresh – casual gaze events that are drawn to it, thus refreshing internal memory.

We saw evidence of this throughout the studies. Documents of importance were placed in obvious locations in full visibility, while less important documents were piled up. There was frequent scanning and refreshing. The location and the visual appearance of documents provided valuable cues that allowed the subjects to remain aware of the state of the document space.

Given the full resolution representations afforded by the space, the memory it provides is almost "perfect", and is, in essence a highly accessibly, on-demand information source. Virtual navigation (e.g., navigating a file system, searching, minimizing, maximizing, etc.) is replaced with *physical navigation* (glancing, turning, moving, etc.) [3]. Location can be considered the "key" to the information.

While it may be faster to perform virtual actions, physical navigation is more efficient because of the embodied resources that support it. There is evidence that the human brain can encode location information alongside inherently non-spatial information such as text [18]. While the evidence for spatial memory is not overwhelming [19], it does not have to be exact. Relocating an object of interest involves both recall and recognition and in a spatial environment, the perceptual system provides many clues, reducing the need for recall. There is also no need to explicitly switch context away from the current task. A glance is enough to grab a piece of information and return to the current task.

The other advantage of full fidelity documents is the integration discussed earlier. For example, one perspective on the role of embedding atomized information as highlights within the document is that the region surrounding the highlighted text is providing rapid access context for the information. Similarly, a document in space can be contextualized by its neighbors.

Semantic Layer

The other key role played by space is that of providing a flexible semantic layer that adds meaning to the displayed information. There are many relationships that can be represented spatially, such as ordering, proximity, and alignment [19]. Space can also serve as a medium for creating complex structures like clusters, lists, and heterogeneous, interrelated types [20]. Using space in this fashion reduces the need for elaborate internal models by replacing memorization and computation with perception. It is, for example, far easier to arrange objects in space and use the perceptual system to *recognize* categories and properties present in a collection of objects than it is to try to memorize all of the characteristics of every object and internally compute relationships.

An advantage of using space to think is that it offers tremendous flexibility, allowing for experimental exploration. Unless explicitly specified (e.g., with a coordinate axis), space imposes no strict interpretation on space. In Robinson's studies of analysts, they adopted a mix of different metaphors as they became useful, sometimes even mixing metaphors in the region of space [8]. The most obvious instance of this in our studies is the clustered timeline discussed previously (Figure 6).

The other benefit of this flexibility is that the meaning of space and organizations can evolve with understanding. At the start of the sensemaking process, when nothing is known about the data, any organizational structure that is chosen is likely to be inappropriate and limiting [10]. By working in a free and completely unstructured environment, the analyst can experiment, making rough categorizations based on impressions or building ad hoc structures that are easily dismantled as understanding changes. As the investigation develops however, understanding starts to grow and with it, a better sense of how to organize the data.

Shipman refers to this use of space as "incremental formalism" and it forms the basis of spatial hypertext [10]. Returning to the timeline example, we find one of the more dramatic instances of this. 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. 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.

We also find evidence for the use of space in the scarcity of notes taken by the subjects. The persistent visual presence of highlighted atoms in the documents and the spatial organization replaced the need to explicitly create further representations.

CONCLUSION

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 this work 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.

While the scenario used in these studies was an intelligence analysis exercise, we believe the results to be broadly applicable. The actual activities that we observed were primarily reading, identifying important information, categorizing, and arranging; all tasks that are common for knowledge workers across many domains.

This work opens a number of interesting paths of future exploration. One of these is the development of longitudinal studies of the use of these displays. Although we observed many spatial uses of the environment by our subjects, it is experts who develop the most revealing strategies for the use of the space [19]. A longitudinal study that examined how domain experts made use of the display after months of use would be very informative.

Another avenue of exploration is to explicitly examine the tradeoffs between physical and virtual space. Many tools (e.g., [9]) have been developed that provide the user with a virtual "canvas" that can be navigated with zoom and pan techniques. How would a virtual space with the equivalent dimensions compare to our purely physical space? Based on previous work and our postulate that the physical environment affords the use of a wider range of human abilities, we would expect the utility of these environments to differ significantly. Given the interest in and accessibility of these virtual spaces, we feel this will be an important area to explore.

At a time when multiple monitor systems are becoming more common, this work makes a case for reconsidering the role of extra display space. We have shown that tiled large, high-resolution displays are not just a collection of monitors – they can work together to become a spatial environment, changing the way the user works and thinks.

ACKNOWLEDGMENTS

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

REFERENCES

1. Pirolli, P. and Card, S. The Sensemaking Process and Leverage Points for Analyst Technology as Identified Through Cognitive Task Analysis *International Conference on Intelligence Analysis*, (2005).
2. Grudin, J. Partitioning digital worlds: focal and peripheral awareness in multiple monitor use. *CHI '01*.ACM Press, (2001) 458-465.
3. Ball, R., North, C. and Bowman, D. A. Move to improve: promoting physical navigation to increase user performance with large displays. *CHI '07*.ACM Press, (2007) 191-200.
4. Czerwinski, M. P., Smith, G., Regan, T., Meyers, B., et al. Toward characterizing the productivity benefits of very large displays. *Interact '03*.IFIP, (2003) 9-16.
5. Shupp, L., Andrews, C., Dickey-Kurdziolek, M., Yost, B. and North, C. Shaping the Display of the Future: The Effects of Display Size and Curvature on User Performance and Insights. *Human-Computer Interaction*, 24, 1-2 (January 2009), 230 - 272.
6. Czerwinski, M., Tan, D. S and Robertson, G. Women take a wider view. *CHI '02*.ACM, (2002) 195-202.
7. Tan, D. S., Gergle, D., Scupelli, P. and Pausch, R. Physically large displays improve performance on spatial tasks. *ACM Trans. Comput.-Hum. Interact.*, 13, 1 (March 2006), 71-99.
8. Robinson, A. C. *Design for Synthesis in Geovisualization*. University Park, PA, 2008.
9. Wright, W., Schroh, D., Proulx, P., Skaburskis, A. and Cort, B. The Sandbox for analysis: concepts and methods. *CHI '06*.ACM, (2006) 801-810.
10. Shipman, I., Frank M., Hsieh, H., Maloor, P. and Moore, J., Michael. The visual knowledge builder: a second generation spatial hypertext. *HYPERTEXT '01*.ACM, (2001) 113-122.
11. Plaisant, C., Grinstein, G., Scholtz, J., Whiting, M., et al. Evaluating Visual Analytics at the 2007 VAST Symposium Contest. *CG & A, IEEE*, 28, 2 (March-April 2008), 12-21.
12. Hutchings, D. R., Smith, G., Meyers, B., Czerwinski, M. and Robertson, G. Display space usage and window management operation comparisons between single monitor and multiple monitor users. *AVI '04*.ACM Press, (2004) 32-39.
13. Hollan, J., Hutchins, E. and Kirsh, D. Distributed cognition: toward a new foundation for human-computer interaction research. *ACM Trans. Comput.-Hum. Interact.*, 7, 2 (June 2000), 174-196.
14. Ramey, T. B. J. Thinking aloud: reconciling theory and practice. *IEEE T-PC*, 43, 3 (Sept. 2000), 261-278.
15. Wise, J. A., Thomas, J. J., Pennock, K., Lantrip, D., et al. Visualizing the non-visual: spatial analysis and interaction with information for text documents. *InfoVis '95*.IEEE, (1995) 442-450.
16. Russell, D. M., Slaney, M., Yan, Q. and Houston, M. Being Literate with Large Document Collections: Observational Studies and Cost Structure Tradeoffs. *HICSS '06*, (2006) 55-55.
17. Heuer, R. J. *Psychology of Intelligence Analysis*. Center for the Study of Intelligence, 1999.
18. Richardson, D. and Spivey, M. Representation, space and Hollywood Squares: Looking at things that aren't there anymore. *Cognition*, 76, 3 (Sept. 2000), 269-295.
19. Kirsh, D. The intelligent use of space. *Artif. Intell.*, 73, 1-2 (Feb. 1995), 31-68.
20. Shipman, I., Frank M., Marshall, C. C. and Moran, T. P. Finding and using implicit structure in human-organized spatial layouts of information. *CHI '95*.ACM Press, (1995) 346-353.