Int. J. Human-Computer Studies 71 (2013) 1078–1088

Contents lists available at ScienceDirect

Int. J. Human-Computer Studies

journal homepage: www.elsevier.com/locate/ijhcs

Large high resolution displays for co-located collaborative sensemaking: Display usage and territoriality

Lauren Bradel [a,*], Alex Endert [a], Kristen Koch [b], Christopher Andrews [a], Chris North [a]

[a] Virginia Tech, 2202 Kraft Drive, Blacksburg, VA 24060, USA
[b] Tulane University, 6823 St. Charles Ave, New Orleans, LA 70118, USA

ARTICLE INFO

Article history:
Received 13 February 2012
Received in revised form
27 July 2013
Accepted 31 July 2013
Communicated by Francoise Detienne
Available online 24 August 2013

Keywords:
Large high-resolution displays
Co-located collaborative sensemaking
Visual analytics
Territoriality

ABSTRACT

From an exploratory user study using a fictional textual intelligence analysis task on a large, high-resolution vertical display, we investigated how pairs of users interact with the display to construct spatial schemas and externalize information, as well as how they establish shared and private territories. We investigated users' space management strategies depending on the design philosophy of the user interface (visualization- or document-centric). We classified the types of territorial behavior exhibited in terms of how the users interacted with information on the display (integrated or independent workspaces). Next, we examined how territorial behavior impacted the common ground between the pairs of users. Finally, we offer design suggestions for building future co-located collaborative visual analytics tools for use on large, high-resolution vertical displays.

1. Introduction

Large displays have been used collaboratively for decades, ranging from digitally enhanced whiteboards to high-resolution computer-driven displays (Rekimoto 1998; Vogt et al., 2011). The physical size of the display, ranging from wall-sized (Beaudouin-Lafon et al., 2012) to desktop-mounted (Endert et al., 2012b), affords multiple users the opportunity to physically navigate the display space instead of users needing to synchronize the virtual navigation of the display. Thus, the large display space allows co-located collaborating users to read at different paces and utilize different analysis strategies, while also sharing their surroundings. These implicit factors such as subtle physical cues, background noises, and additional contextual items contribute to common ground (Chuah, 2003).

The use of large displays for collaborative analysis of free-form text documents has expanded in recent years, particularly with tabletop and projector-based displays (Birnholtz et al., 2007; Isenberg et al., 2010). Large, high-resolution displays (LHRDs) composed of vertical LCD displays expand upon the large display surface by easily displaying whole documents at standard magnification levels. This allows users to place detailed views of documents (as opposed to thumbnails or labels) into spatially meaningful representations, which can then be used to easily recall information through physical navigation, as well as to semantically organize the display space (Andrews et al., 2010). These properties of large, high-resolution displays have been shown to improve user performance on many tasks ranging in difficulty from simple pattern finding and route tracing to cognitively demanding sensemaking (Andrews et al., 2010; Ball et al., 2007). Visual analytics tools designed to aid the sensemaking process can be designed to leverage the spatial affordances of the display (Endert et al., 2012a).

Collaborative visual analytics has been a growing research area within the visual analytics community due to the ability to integrate social and group dynamics into the analytic process (Heer, 2008; Thomas and Cook, 2005). The intersection of pair collaboration on sensemaking tasks and large, high-resolution vertical displays presents an opportunity for specially designed visual analytics tools. However, designing collaborative tools specifically for large, high-resolution vertical displays requires a better understanding of user behavior in terms of territoriality and display space management.

We seek to remedy this knowledge gap by analyzing data from an exploratory study of two different tool design philosophies that differ in how they use the display space: visualization-centric and document-centric design. Both use multiple windows to display more information on the large display. The difference lies in what the windows are used for. Visualization-centric tools emphasize the use of multiple coordinated visualizations of document meta-data such as extracted entities in the dataset, and include a tabbed document browser for viewing detailed text content. This design tends toward a static arrangement of a fixed set of views. Visualization-centric tools

* Correspondence to: 14107 Stroubles Creek Rd., Blacksburg, VA 24060, USA. Tel.: +1 412 874 7807; fax: +1 540 231 9218.
E-mail addresses: lbradel1@cs.vt.edu (L. Bradel), kkoch@tulane.edu (K. Koch), cpa@cs.vt.edu (C. Andrews), north@cs.vt.edu (C. North).

1071-5819/$ - see front matter © 2013 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.ijhcs.2013.07.004

often provide a designated synthesis space (e.g. Jigsaw's Tablet view (Stasko et al., 2008)), but otherwise focus on visualizations for foraging. Coordination between views is often accomplished via brushing and linking. One example of a tool embodying this visualization-centric design philosophy is the visual analytics tool Jigsaw (Stasko et al., 2008).

Document-centric tool design, on the other hand, emphasizes the use of multiple document viewers (instead of iconic visualizations) in multiple separate windows with one document per view. Data is displayed in its natural form of unstructured text, perhaps with entity extraction or highlighting, and offering manual text mark-up. This design philosophy encourages a dynamic, user driven, and unrestricted layout arrangement, and often attempts to mimic how users might organize paper documents on a large table as a cognitive synthesis space. One example of a tool embodying this document-centric design philosophy is Analyst's Workspace (AW) (Andrews and North, 2012).

These categories are not intended to label all possible analytic tools. Rather, they represent two different approaches for conveying information and managing available display space. They are not entirely mutually exclusive either, but represent emphases in tool designs.

Previously, we investigated analysis strategies and user roles that developed during a two-hour text-only sensemaking task (Vogt et al., 2011). In this paper, we detail how visual analytics tool designs afford different spatial usages of large, high-resolution vertical displays and which collaborative styles they tend to encourage users to adopt. We found that the document-centric tool users referenced more information by spatial location and schematized the display via spatial metaphors than the visualization-centric tool users. We also present our findings on the impacts of territorial behavior and collaboration styles on the effectiveness of collaborative sensemaking. In terms of territoriality, we found that teams that adopted an integrated workspace (large shared territory) communicated more and ultimately had more similar solutions than teams that adopted an independent workspace (large personal territories). Finally, we conclude with a discussion on how these results translate into design suggestions for future collaborative visual analytics tools for text data on large, high-resolution vertical displays.

2. Related work

Co-located collaborative sensemaking has been studied in various domains (Morris et al., 2010; Paul and Morris, 2009; Paul and Reddy, 2010). However, user requirements vary across domains due to the specific nature of the work. For example, the competitive workplace culture of intelligence analysts means that collaboration occurs only informally, if at all (George Chin et al., 2009; Heuer and Pherson 2010). Therefore, it is important to design tools and environments with little overhead required to commence collaboration (Gutwin, 2008). This fact motivated us to present users with existing tools not designed specifically for collaboration with minimal system changes, similar to the retrofitting of NodeTrix for co-located collaboration, resulting in a tool called CoCoNutTrix (Isenberg et al., 2009). Their analytic environment consisted of different configurations of large, projector-based vertical displays, multiple independent mice, and a single keyboard. They added collaborative functionality to NodeTrix as simply as possible, using off-the-shelf products and existing toolkits when possible. No conflict resolution support was provided, allowing the participants to verbally resolve conflicts. The participants overwhelmingly did not feel that they needed such support. Because participants could not open separate instances of a NodeTrix window, as they could in our study, it is not surprising that they chose to resolve conflicts through communication.

Although pair dynamics exist, and are generally well understood, between domain experts and tool experts (Green et al., 2009; Pickens et al., 2008), we sought to better understand the collaborative process between equally knowledgeable collaborators, such as between co-workers working on a joint investigation. Software that can support this collaborative process, such as Single Display Groupware (SDG) (Stewart 1997), has been studied extensively in the past, starting with early systems in the late 1980s and early '90s (Elrod et al., 1992; Pedersen et al., 1993; Rekimoto 1998). In subsequent work ((Stewart et al., 1999; Stewart et al., 1998)), Stewart et al. investigated SDG systems further. They then conjectured that "very limited screen space […] may result in reduced functionality compared with similar single-user programs" (Stewart et al., 1999). Increasing the display screen's physical size, and subsequently resolution, to provide adequate virtual and physical space for SDG systems can alleviate this concern.

Design decisions that enhance individual work can often hinder group work, and vice versa. Previous groupware interfaces have either supported group work through consistent view sharing, known as "What You See Is What I See (WYSIWIS)," or the individual user through more relaxed view sharing (Stefik et al., 1987). As Gutwin and Greenberg state, "the ideal solution would be to support both needs – show everyone the same objects, as in WYSIWIS systems, but also let people move freely around the workspace, as in relaxed-WYSIWIS groupware" (Gutwin and Greenberg, 1998). We believe that a suitable balance can be reached between these tensions by allowing users to work on a large, high-resolution display where they can work individually while maintaining the ability to see their collaborator's actions.

Large displays come in many different form factors. These include, but are not limited to, LCD or projector displays, and vertical or horizontal displays. One factor in deciding which display to use is whether or not the display provides adequate space for personal, shared, and storage territories to form as the participants see fit (Scott, 2004). Territoriality and other co-located design issues have been studied on tabletop displays using systems such as Lark or Cambiera (Isenberg et al., 2008). The tabletop display used to study the Cambiera system was 2×3 in^2. with a resolution of 1024×768 pixels, while Lark (Tobiasz et al., 2009) was run on a larger 5×3.28 in^2. display with a resolution of 2800×2100 pixels. Some study participants on the display used with Cambiera commented that they "felt cramped, wanting a higher-resolution and physically larger display for document reading" (Isenberg et al., 2010). Thus that particular physical set-up did not provide "appropriate table space" to define territories (Scott, 2004). Note that these observations were not present with the Lark system.

Large vertical displays can range from wall-sized projector displays (15×5.4 in^2. in size with a resolution of 14 megapixels) to desktop LCD displays (4×2 grid of 30 in. LCD monitors with a total resolution of 32 megapixels) (Andrews et al., 2010; Birnholtz et al., 2007). Both of these set-ups are large enough to support the physical space Scott et al. claims supports the development of territories (Scott, 2004). Although the projector display is much larger than the LCD display, its resolution is much lower. Therefore, the higher resolution LCD display is better suited to close-proximity document viewing. We chose to investigate large, high-resolution vertical displays composed of LCD displays for the purpose of investigating territoriality during a text-heavy analytical task. We felt that this type of display would provide a physical environment that allowed for territories to develop while maintaining a high resolution for text reading.

Given the choice to use a large, high-resolution vertical LCD display, the next decision to make is the number of input devices (mice and keyboards) to use. Stewart et al. found that one device per person was preferable in SDG systems because they increased interaction and kept both participants "in the zone" (Stewart et al., 1998). Although it has been shown that multiple input devices allow for more parallel work but less communication (Birnholtz et al., 2007), multiple input devices allow for more reticent participants to contribute to the task (Robinson, 2008). Because we sought to keep users in the "cognitive zone" (Green et al., 2009) and allow for simultaneous text input as well as mouse input (Isenberg et al., 2009), we chose to implement two mice and keyboards, one for each user, to enable them to contribute to the collaborative sensemaking task simultaneously instead of frequently trading mouse and keyboard ownership.

3. Study description

We conducted an exploratory user study to observe the co-located collaborative sensemaking process with text documents on a large, high-resolution vertical desktop display.

In order to observe a wider range of user behavior, we chose to study two tools that came from different tool design philosophies: visualization-centric and document-centric. We chose these two contrasting tools to analyze how users would manage the large, high-resolution display space in terms of organization and territoriality. From the visualization-centric design philosophy, we chose Jigsaw (Stasko et al., 2008), which uses multiple coordinated views of entity-centered visualizations as well as a designated synthesis space. From the document-centric design philosophy, we chose AbiWord (www.abisource.com), which is a simple document viewer that allows manual user highlighting in a variety of colors and a simple file browser tool to allow searching across documents. We chose these tools because they each represent their corresponding design philosophy in the extreme. We hypothesized that the document-centric tool (AbiWord), with flexible layout of multiple AbiWord instances offered by the window manager, would better take advantage of the display space to promote more collaborative activity than the visualization-centric tool (Jigsaw), due to the ability of LHRDs to provide detailed views of textual information while maintaining spatial representations, even though visualization-centric tools have performed well in single-user studies (Andrews et al., 2010; Kang et al., 2009; Stasko et al., 2008).

Additionally, we allowed territories to develop without restriction by providing two mice and two keyboards (one per user) that could operate independently with two distinct on-screen cursors. Using two mice and keyboards allowed the users to choose how much of their time was spent working together or separately. Each mouse could simultaneously maintain its own active window [Fig. 1].

3.1. Research questions

In the context of two users performing a text analytics sensemaking task using one of two types of tools (visualization-centric or document-centric) on a shared large, high-resolution display: How do the users organize and manage information on the display? How do the users collaborate and develop territories? How do these behaviors impact sensemaking?

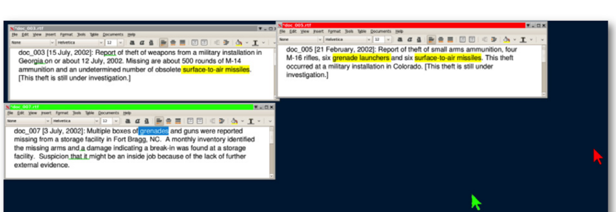

Fig. 1. Two mice with two active windows.

3.2. Participants

We recruited eight pairs of participants (J1–J4 used Jigsaw; D1–D4 used the document viewer). Six of the eight pairs were students and the other two pairs consisted of research associates and faculty, and all pairs knew each other prior to the study and had previous experience working collaboratively. There were four all male groups, one all female, and three mixed gender. Each participant was compensated $15 for participation. As a form of motivation, the verbal debriefing solutions formed by the teams of participants were scored and the participants received an additional financial award for the four highest team scores.

3.3. Workspace set-up

The teams of users sat in front of a 108.5 in. × 35 in. display consisting of a 4 × 2 grid of 30-in. LCD 2560 × 1600 pixel monitors totaling 10,240 × 3200 pixels, or 32 megapixels [Fig. 2]. The display was slightly curved around the users, letting them view the majority, if not all, of the display in their peripheral vision. A single machine running Fedora 8 drove the display. A multi-cursor window manager based on modified versions of IceWM and x2x was used to support two independent mice and keyboards (Vogt et al., 2011). Thus, each user was able to type and use the mouse simultaneously and independently in the shared workspace [Fig. 1]. A whiteboard, dry-erase markers, paper, and pens were also available for use because these external artifacts were explicitly requested during the pilot study. The whiteboard was located approximately four feet behind the seated participants. Each participant was provided with a rolling chair and freestanding, rolling tabletop holding the keyboard and mouse so that they could move around if they chose to do so. The desks and chairs were positioned side-by-side in front of the display.

3.4. Analytic tools

As mentioned previously, we chose to investigate two different types of analytic tools. This is by no means an exhaustive survey of visual analytic tools, but this design decision allowed us to observe a wider range of display usage and collaborative behavior than if we had only observed a single type of tool.

3.4.1. Jigsaw

Jigsaw (Kang et al., 2009; Stasko et al., 2008) is a text analytics system that has been designed to support analysts in the sensemaking

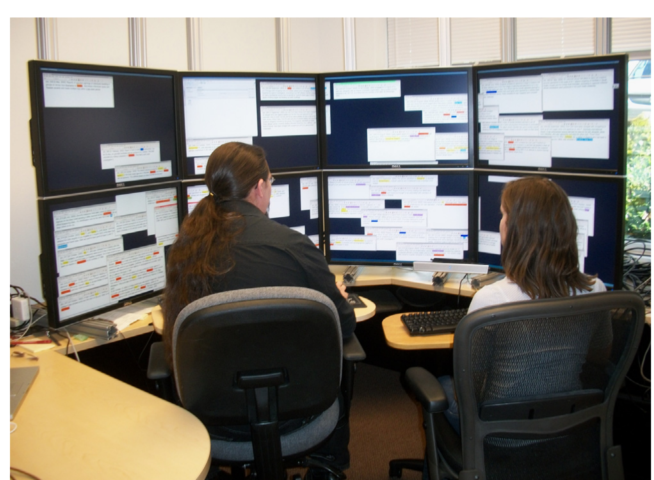

Fig. 2. Various Jigsaw views. From left to right, top to bottom: list view, document view, tablet view, graph view, document view, Jigsaw startup window.

process, and has proven to be an effective tool in multiple user studies and case studies. Jigsaw visualizes document collections in multiple views based on the entities (people, locations, etc.) contained within those documents [Fig. 3]. All entities are highlighted and color-coded automatically when Jigsaw imports a document collection. It also allows textual search queries of the documents and entities. Jigsaw can sort documents based on entity frequency, type, and relations, and this information can be displayed in a variety of ways, including interactive graphs, lists, word clouds, and timelines. These multiple views are linked to each other by default so that a selection or change in one window propagates to the remaining views. Jigsaw also provides a Tablet view where users can write notes, draw connections between entities, and create timelines. Because of its use of multiple coordinated views, Jigsaw is designed to work well in multiple monitor configurations. When users want to view the contents of multiple text documents, Jigsaw′s default behavior is open the documents in a single tabbed browser. Users flip through the documents by selecting the tabs. While users are able to open multiple such browser instances, all browsers will display the same documents unless users execute a series of non-trivial steps. Thus, Jigsaw tends to devote screen space to data visualizations rather than document views, making Jigsaw a *visualization-centric* tool. Because of the complexity of Jigsaw, participants were given a tutorial prior to the start of the task.

Jigsaw was not designed as a collaborative tool. Rather than retrofitting the tool (such as in (Isenberg et al., 2009)), we gave the participants a single instance of the tool in order to facilitate awareness regarding the other user′s actions (such as in (Isenberg et al., 2010)). Users were able to open multiple instances of the same view. However, Jigsaw′s brushing and linking feature acted as if only one user was operating the system instead of tracking two separate threads of interaction. Uninterrupted parallel work could be conducted if one or both participants turned off this feature.

3.4.2. AbiWord (document viewer)

The second tool that we chose was a basic document viewer, AbiWord, which allows users to display documents, manually highlight document portions, edit existing documents, and create text notes [Fig. 4]. Teams using this document viewer were also provided with a file browser tool in which they could search for keywords across the document collection and open documents. Opening a new document resulted in a new AbiWord instance window on the screen, prompting users to spread documents out on the display using the standard window manager functionality. Hence, we classified this tool as *document-centric*, because it devoted screen space almost entirely to viewing document contents with the only exception being the file

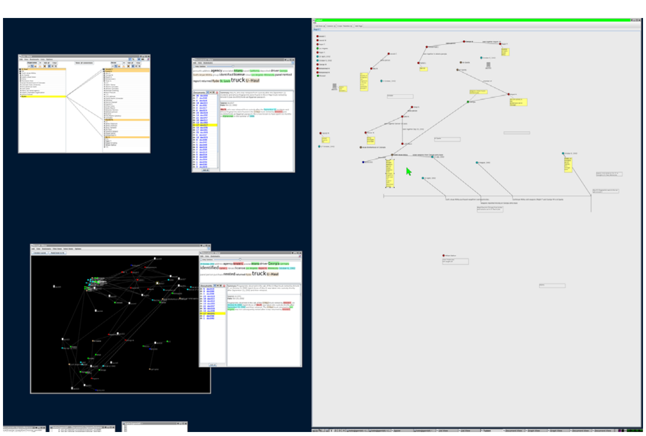

Fig. 3. Two users seated in front of the large, high-resolution set-up equipped with two mice and keyboards used in the study.

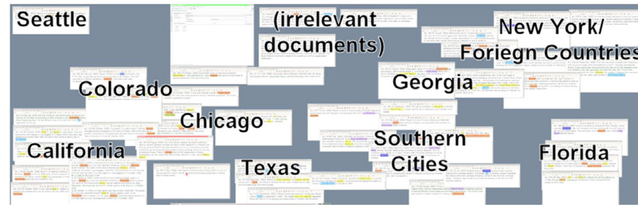

Fig. 4. Geographical document clustering.

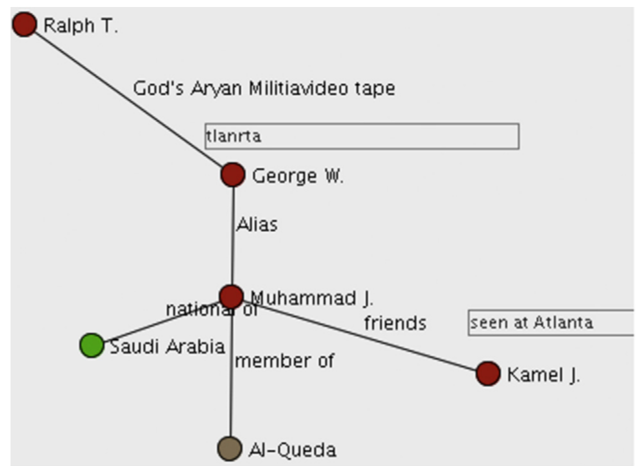

Fig. 5. J4 zoomed in Tablet view showing connections between entities.

browser window(s). Participants were given a brief overview of this tool.

A file browser window was provided to the participants, which contained all text documents for the study in a single directory. Searches were therefore done in series. Parallel, or simultaneous, searching could be done by duplicating the file browser window.

3.5. Task and procedure

After the tutorials on Jigsaw or the document viewer with a sample set of documents, each team was given two hours to analyze a set of 50 text documents and use the information gathered to predict a future terrorist attack on the United States. The scenario used in this study comes from an exercise developed to train intelligence analysts and consists of a number of synthetic intelligence reports concerning various incidents around the United States, some of which can be connected to gain insight into a potential terrorist attack. This same scenario was also used in a previous study evaluating individual analysts with Jigsaw (Kang et al., 2009).

3.6. Data collection

Following the completion of the task, each participant filled out a report sheet to quantitatively assess their individual understanding of the analysis scenario, then verbally reported their final solution together to the observers. The rubric for evaluating the participants′ verbal and written solutions was based on the strategy for scoring the Visual Analytics Science and Technology (VAST) challenges (Plaisant et al., 2007). The participants earned positive points for the people, events, and locations related to the solution and negative points for those that were irrelevant or incorrect. They also received points based on the accuracy of their overall prediction of an attack, guidelines for which were provided in the scoring rubric. Negative scores were possible for the participant′s total score. The joint verbal debriefing was scored to produce the group′s overall score. There were ten possible points for person identification, fifteen possible points for important event identification, four possible points for location identification, and seven possible points for the overall prediction. The highest score a team could get for this dataset was 36 points.

The individual reports filled out by the participants were compared against their teammate′s to calculate similarities and differences between their solutions. This similarity score was calculated by granting one point for each common entity listed or mentioned in both individuals' reports, and one point was taken away for each entity that was present in only one individual′s written report. Thus, negative scores were possible for this metric.

Additionally, individual semi-structured interviews were conducted where each participant commented on how they solved the scenario, how they arranged information on the display, and how they felt collaboration affected their ability to solve the scenario.

During each study session, an observer was present taking notes. Video and audio of every scenario, debriefing, and interview was recorded. We also automatically collected screenshots in fifteen second intervals, logged mouse actions (movements and clicks), and logged active windows.

3.7. Data analysis

We used the overall solution correctness scores to identify any significant differences based on which analytic tool was used. This relationship was not significant, although the sample size was small, making significance difficult to accurately determine.

Additionally, we compared the individual solution reports generated by each pair′s participants to calculate their shared knowledge as it pertained to the solution of the scenario, calculating the amount of common ground in the solution.

The screenshots were combined with answers from the semi-structured interviews to identify what organizational strategies were used on the display, if any. We also referenced the screenshots to observe the use of display space throughout the study and to “play back” the study session along with the recorded video due to resolution limitations of the video. Finally, the screenshots were used to calculate the amount of empty space, or whitespace, on the display to quantify the amount of the display that the pairs used.

The videos were coded to calculate the percentage of time the pairs spent closely collaborating during the session, using a coding set established by Isenberg et al., 2009 where close collaboration involves active discussion, working with the same documents, or working on the same specific problem, and loose collaboration involves working on similar problems from different starting points, different problems, or one participant is disengaged from the task (Isenberg et al., 2010). We also transcribed the video in order to quantify verbal cues that were linked with different types of territorial behavior.

We used the mouse data to further establish territorial boundaries that existed between the participants, if any. We accomplished this by totaling the number of mouse button-down events that occurred in each of the eight screens of the display for each participant. Using this information, we could see which screens were used by one participant, both, or none. The mouse data was also used to calculate the percentage of clicks in each display screen to identify where each participant primarily interacted with the display.

4. Information organization and management

To answer our first research question (*How do the users organize and manage information on the display*?), we analyzed how the groups, partitioned by analytic tool, used the large display space to externalize information to aid their sensemaking process. We hypothesized that the document-centric tool (document viewer) groups would use a higher percentage of screen space and that the display space would be used in a dynamic way to represent semantics in their findings (similar to (Andrews et al., 2010)). We hypothesized that the visualization-centric tool (Jigsaw) groups would use a lower percentage of the display space and impart less meaning upon the space as a static spatial layout, favoring the use of visualizations and Jigsaw′s designated synthesis window. In order to investigate display space management and organization, we referenced the user interviews explaining their organizational schemas (section 4.1) and utterances regarding re-finding information previously externalized to the display (section 4.2).

4.1. Analytic actions and display usage

Two strikingly distinct approaches were taken towards utilizing the display space between the study conditions (visualization-centric or document-centric).

As suspected, the visualization-centric groups favored leveraging visualizations of groups of entities to conduct their analysis. These groups (J1–J4) did not use the entire display space. After opening and arranging various visual analytics views in Jigsaw, these windows rarely changed position on the display. Participants in these groups only opened one or two documents at a time in Jigsaw′s document viewer. Three out of four Jigsaw groups used Jigsaw′s Tablet view to record connections between people, places, and events, while the fourth team used paper to accomplish this. The groups that chose to use the Tablet view spatially arranged entity information in this virtual space.

For the document-centric groups, up to 50 documents were simultaneously displayed, which is not surprising because the file browser opened each document in a new window, as opposed to Jigsaw′s Document View, which replaced each open document with a new one. These groups (D1–D4) did not have access to advanced features, such as connecting entities across documents. Instead, their only method of learning the contents of the document collection was to read every document. All document viewer groups chose to record notes on paper or a whiteboard. No groups recorded notes on the computer. While reading the documents, all document viewer groups highlighted seemingly relevant entities and arranged the documents on the display. All groups initially opened all 50 documents in separate windows on the display, and only closed documents once they were deemed irrelevant to the solution. Entity highlighting, which the groups had to do manually, was done by all groups even though the proctor did not guide the groups on how to conduct their investigation. The arrangement of documents fell into several organizational schemas based on content, temporal, or geographical relevance.

These spatial arrangements differed from spatial arrangements constructed by groups using the visualization-centric tool′s dedicated synthesis space. We sought to further analyze these differences.

4.1.1. AbiWord clusters

All document-centric groups clustered the documents using an overall organizational scheme. Two groups formed clusters based on content similarity (e.g. people, organizations, events occurring in multiple documents). One group organized their display space geographically based on locations of events in the documents, mentally superimposing a map of the United States on the large display, with foreign countries located where the Atlantic Ocean would be. The final group created multiple timelines on the display to track the evolution of events or track individual′s actions. All clustering schemas were established through discussion between group members early in the investigation. No clusters were explicitly labeled on the display.

Clusters were generally identified on the display by adjacent or overlapping documents, and were separated from different clusters by whitespace (i.e. empty space showing just the desktop background) on the display. In this manner, the document viewer groups transformed the "unused" portions of the display to aid their cognitive process regarding accessing information externalized onto the display.

The clusters formed by the document viewer groups expanded across the entire display [Fig. 4]. Because of the large display space, participants were able to display all 50 documents simultaneously, eliminating the need to use icons. This strategy of keeping all documents visible may not be possible with extremely large document sets. Because participants freely overlapped documents with this small dataset, we believe it is reasonable to extrapolate that users completing an analytical investigation using a document-centric tool for a large dataset would still cluster documents. These clusters would likely resemble piles with a greater amount of document overlap and that these clusters would still carry spatial meaning.

4.1.2. Jigsaw tablet-view clusters

Groups in the visualization-centric study condition did not cluster windows. Instead, the Jigsaw groups that chose to use the Tablet view formed a different type of clusters that were composed of entities, not entire documents, and were contained in the Tablet view, as opposed to expanding across the entire display. Multiple clusters of nodes (entities) were created, in contrast to clusters composed of documents.

As opposed to generic relationships such as "Chicago" or "U-Haul," the visualization-centric groups drew explicit connections based on the contents of clusters which contained labeled edges that depicted the relationship between nodes, such as:

"*Muhammad J., who is an alias for George W., is a member of Al-Queda and is friends with Kamel J.*" [Fig. 5]

This was a much more formal method of clustering than was seen in the document viewer groups due to the explicitly labeled connections between nodes.

It should be noted that the Jigsaw groups did not utilize their whitespace as the document viewer groups did. They did not spatially organize the different Jigsaw views into any meaningful arrangements, as evidenced through interview questions regarding their display usage. For the Jigsaw groups, the whitespace between views was merely empty and unused space. These groups, however, did use whitespace in the same manner as the document viewer groups within the Tablet view to separate clusters of entities.

4.2. Externalization for Re-finding information

All teams externalized information, although the methods of externalization and what was externalized differed based on the analytic tool used. In addition to notes explicitly written by the participants on paper or on the display, users maintained information on the display in order to reference it later in their investigation. This process of saving information for later reference was done differently with each tool.

In the visualization-centric tool (Jigsaw), the Document Viewer window automatically replaced previously viewed documents with the current document selected. It is not surprising that users did not go through the extra effort required to open documents in separate views in order to construct a spatially meaningful representation across the large display that could be referenced by physical navigation. Jigsaw groups typically revisited documents through virtual navigation (e.g. selecting a tab). Aside from documents, separate views were spread across the display, although this configuration was primarily static and did not evolve from their initial configuration. The actual configuration carried little semantic meaning for the groups using Jigsaw, evidenced through the post-study interviews, where the document-centric tool groups described the layout of the display and the Jigsaw groups made no mention of it.

To reference information previously read, participants from the visualization-centric group utilized the dedicated synthesis space provided in the tool (Jigsaw′s Tablet view). Participants recorded notes consisting of connections between entities or events that the participants had established. These included the travel history of individuals, lists of suspects, their aliases and allegiances, as well as notable or suspicious events. This information was sometimes contained on a timeline, a feature available in Jigsaw′s Tablet view. Overall, these notes contributed to hypotheses concerning the fictional terrorist plot the participants were attempting to uncover.

In the document-centric tool (AbiWord), the file browser opened documents in separate windows. It is unsurprising, given this behavior, that participants kept a separate window open for each document. These participants arranged the text documents on the display in spatial schemas. In addition to this being an organizational strategy, as described above, users were frequently able to re-find information contained in documents by spatial location and completed few searches on the computer by comparison. Many locations of documents in the spatial layouts were explicitly discussed, making them physical representations of shared knowledge, similar to the evolving workspace serving as an archive of analysis in the CoCoNutTrix study (Isenberg et al., 2009).

In terms of externalizing synthesis-related thoughts, the document-centric groups wrote on physical paper or the whiteboard close behind them (no virtual synthesis space was provided as in the other study condition). These notes consisted either of document numbers to keep track of which were read (Jigsaw indicates this automatically), or entities written down, as seen in the visualization-centric groups.

All pairs of collaborators, regardless of analytical tool used, used this large, high-resolution display as a form of external memory, as shown by their references to the space itself. All groups referenced documents by their spatial location and contents more than by document title or number. For example, participants used words such as "here" or "there" to indicate position, often accompanied by a pointing gesture towards the indicated region of the display:

"*There are some surface-to-air missiles up there.*"

Occasionally, documents were referenced by their name:

"*[Document] 35 is just right above it.*"

This pointing behavior, accompanied by utterances regarding physical location, appeared to indicate that spatial position was important for participants to re-find information. Participants using the document-centric tool pointed significantly more at the screen than participants using Jigsaw [Table 1], apparently as an efficient form of communication for collaboration. Additionally, all groups used the physical location of information on the screen to re-find documents or information [Table 1]. In Jigsaw, this was typically done through the graph, list, or tablet views. All but one group re-located information by using spatial references more often than using a search function on the computer to locate the entity or document in question.

5. Territoriality

To answer our second research question (*How do the users collaborate and develop territories, and how do these behaviors impact sensemaking*?), we analyzed the territoriality and collaboration styles of the pairs, partitioned by groups that viewed the display as an entirely shared space and those that viewed it as containing shared

Table 1
Partitioned by tool used (J: Jigsaw; D: Document Viewer), scores calculated using the VAST challenge rubric (from verbal debriefing), score similarity (from individual written reports), average amount of whitespace on the display, total number of notes taken, number of times participants pointed at the display, number of times participants re-found information through computerized searches or by spatial reference, and number of times participants referred to documents by the document name or by spatial location.

Group	Total score	Report similarity	Average % whitespace	Notes taken	Pointing count	Re-find by search	Re-find spatially	Doc. refer by name	Doc. refer spatially
J1	11	8	86.77%	140	76	4	11	4	16
J2	−1	4	55.60%	90	43	6	7	2	20
J3	−2	3	86.84%	121	122	3	2	2	5
J4	−7	−17	27.24%	91	97	0	7	2	19
D1	13	2	61.23%	44	211	8	27	11	74
D2	−1	−26	50.88%	153	95	1	12	4	33
D3	10	4	54.80%	24	115	0	3	1	27
D4	14	10	51.64%	147	165	4	11	3	44

and individual partitions. We chose this partition because we wanted to investigate how users establish territories on a large, high-resolution vertical desktop display, regardless of tool type used (granted, this is not an exhaustive survey).

We hypothesized that a closely coupled collaboration style would be linked with better performance on the intelligence analysis task and that this closely coupled style would be linked with primarily shared display space, as opposed to the use of large private or personal territories.

5.1. Territorial style characteristics

We classified the groups into two collaboration styles based on how they shared (or did not share) the display space. These styles were Integrated Workspace [Fig. 7] and Independent Workspace [Fig. 6]. The characteristics of these styles can be seen in [Table 2]. The classification of groups was based on the dialog between participants, how they transferred documents across the display, mouse click distributions, and video coding of the closeness of collaboration using an established set of codes (Isenberg et al., 2010).

Integrated groups tended to use plural possessive pronouns such as “our” in speech, such as “we know that…” even if only one person had read the document containing the referenced information. The independent groups, on the other hand, used singular possessive pronouns, such as “my” or “your” to reference the same kind of information. For example (emphasis added):

Independent Workspace: *“You are not stealing **my** document!”*

Integrated Workspace: *“I like these three people as **our** suspects.”*

As seen in [Table 3], groups that used the large, high-resolution display as an integrated workspace tended to collaborate more closely and score higher on the intelligence analysis scenario than their independent workspace counterparts.

Typically, integrated workspaces were associated with the document-centric tool groups, while the independent workspaces were associated with the visualization-centric tool groups, but this was not always the case. This behavior could be attributed to the affordances of each tool. AbiWord produced document windows that quickly filled the space, and duplicated documents may have grown confusing and difficult to keep straight. On the other hand, Jigsaw did not fill up much space, leaving open the possibility of conducting analysis on a large number of documents or entities in a smaller amount of space. Granted, neither of these factors forced users into a specific territorial behavior. When compared to the CoCoNutTrix system, we notice that the lack of conflict resolution support may have led to independent workspaces in the Jigsaw groups (Isenberg et al., 2009). Instead of working in the same analytic workspace and mitigating conflicts, participants could simply open a second instance of a Jigsaw view.

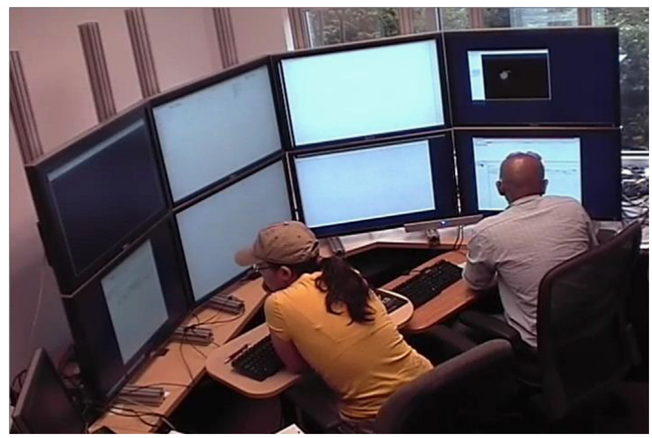

Fig. 6. An independent workspace group working on separate threads of the investigation

Fig. 7. An integrated workspace group working to solve a joint problem.

As seen in [Table 3], the integrated and independent workspace groups are 3/4 composed of one analytic tool (AbiWord for integrated, Jigsaw for independent). However, two groups, J1 and D2, prove the exception to this trend, illustrating that the analytical tool used did not restrict the collaboration style adopted by the participants. The overall trend appears to suggest that this document-centric tool affords an integrated workspace, while Jigsaw affords an independent workspace. However, this trend could be attributed to usability restrictions of the tools.

Table 2
Characteristics of integrated and independent workspaces.

elsels	Independent workspaces
Few apologies, if any	Apologies when other's workspace is "invaded" (indicates "personal space")
Information passed freely across the screen	Information moved to the central shared space when sharing it
Plural possessive pronouns ("our")	Singular possessive pronouns ("my" and "your")
Mouse clicks more evenly distributed across the display	Mouse clicks biased towards each person's side of the display

Table 3
Partitioned by collaboration style, overall score earned, similarity of individual reports, and percentage of time spent in close collaboration.

Group	Collab. style	Total score	Report similarity	% Close collab.
J1	Integrated	11	8	93.39%
D1	Integrated	13	2	98.11%
D3	Integrated	10	4	89.69%
D4	Integrated	14	10	98.78%
J2	Independent	−1	4	45.82%
J3	Independent	−2	3	67.24%
J4	Independent	−7	−17	42.04%
D2	Independent	−1	−26	54.75%

For example, although Jigsaw's tablet view could be resized and did not restrict users on information placement, two users could not type in it simultaneously, which forced users to take turns or open separate Tablet views. Out of the three groups that used the Tablet view, only one opened two separate Tablet views, while groups J1 and J2 both chose to maintain one joint Tablet view. In group J1, one of the outlier groups that adopted an integrated workspace, one participant primarily added notes to the Tablet, which prompted conversation regarding the information entered. Group J1 was able to use their shared Tablet successfully, primarily because they maintained high levels of communication, including what information to enter into the Tablet.

Group J2 (an independent workspace group), however, encountered more difficulties when attempting to use a shared Tablet. They sometimes overrode the other person's actions in the shared view, resulting in apologies between the pair. These types of overriding actions were present throughout group J2's investigation:

> Participant B: *"You are not stealing my document!"*
>
> Participant A: *"Yes I am."*
>
> Participant B: *"Wait, I need this document. Why don't you just copy it?"*

In contrast, group J1 was able to work together using shared Jigsaw views. This required a high level of communication between this pair that was not exhibited by the other teams that used Jigsaw. The consistently close collaboration exhibited by group J1 resulted in a shared, integrated workspace, and also a higher score than any other Jigsaw group [Table 2, Table 3].

The other outlier group, D2, experienced conflicts that arose from a lack of communication. Because they infrequently communicated their thoughts and actions, they often intruded on the other person's actions without realizing they were:

> Participant B: *"Hey, hey, hey – watch it! I'm highlighting. What the heck are you doing?"*
>
> Participant A: *"I'm just trying to move it."*

These types of interactions occurred throughout the course of their investigation. The conflicts between the members of group D2 may have prompted them to adopt individual territories, or the individual territories may have prompted the conflicts.

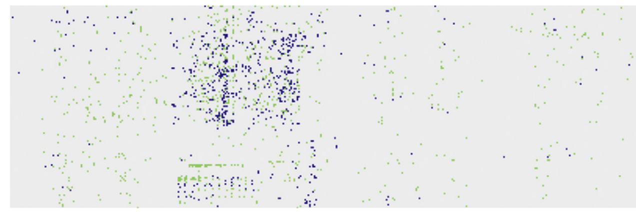

Fig. 8. Group J2: mouse clicks with clear boundaries established.

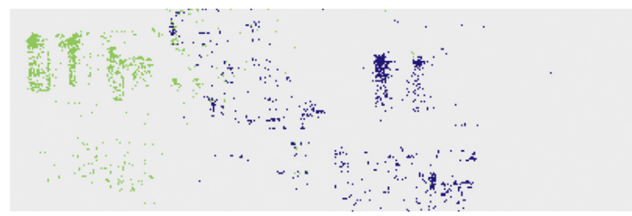

Fig. 9. Group D1: mouse clicks distributed across the display with no clear boundaries.

Overall, group D2 only collaborated closely 54.75% of the time spent solving the scenario and produced the most dissimilar individual reports with a similarity score of −26 [Table 3]. The dissimilarities of their individual reports suggested that they found different entities and hypotheses to be more important and relevant to the correct solution. Their group verbal debriefing was marked by turn taking where each individual filled in gaps of the other person's story. Therefore, together they were able to present a solution whose correctness score was tied for fifth highest out of the eight teams. Thus, they did not perform terribly, but they also did not excel at the intelligence analysis scenario.

These outliers demonstrate that the type of tool used does not strictly dictate collaborative strategy and territorial behavior, even though the design of the tool may better afford one type of collaborative and territorial style. This begs the question of how to design collaborative analytical tools specifically for use on large, high-resolution displays that facilitate close collaboration.

5.2. Shared display space

For those groups that adopted an integrated workspace, the mouse clicks were evenly distributed across the display without any "off-limits" areas for individual space [Fig. 9], while the independent workspace groups showed partitions of shared and individual territories [Fig. 8]. Even though boundaries were never formally vocalized, they could still be detected by analyzing the mouse data collected. The individual territories formed in front of each participant, while the shared territories formed in between the participants, as in (Scott 2004). However, no storage territories developed. This could be attributed to the size of the document collection, but is worth noting. There was a bias toward integrated and independent workspaces for document-centric and visualization-centric tools, respectively, but this was not always the case.

Interestingly, once the independent workspace groups subconsciously established individual and shared boundaries, they apologized whenever they inadvertently crossed the imaginary lines they had drawn or interfered with the other person's actions.

"Sorry, did Georgia display in your window?"

Although rare, there were instances where participants explicitly chastised their partner for invading their personal territory and the documents contained in it:

"Why are you looking at my documents?!"

In order to pass documents or windows across the display when shared and individual territories were established in the independent workspace groups, participants dragged windows to the shared territory and allowed the other participant to drag the window from the shared territory to their own individual area. These groups did not extensively share documents. In integrated workspace groups, participants did not hesitate to drag documents or windows across the entire display, unhampered by any invisible barrier. Storage territories consisting of clusters or piles commonly referred to as "irrelevant documents" were found in a few of the teams (document-centric tool users only). These piles were located on the perimeter of the display and were treated similarly to clusters throughout the display.

It is not surprising that the amount of shared display space is closely linked to the amount of close collaboration between the participants. Without individual territories in which the participants could work on separate threads of the investigation, they tended to work together to solve the scenario. The independent workspace groups, on the other hand, often secluded themselves on "their side" of the display to pursue hypotheses, eventually coming back together to discuss their findings. The territories afforded different styles of collaboration. The closely collaborating groups scored much higher on the scenario, and also reported more similar solutions than those groups that collaborated loosely. This phenomenon is tightly linked with how successful the participants were at attempting to achieve and maintain common ground.

5.3. User roles

Five out of eight groups developed clearly defined user roles, which we previously classified as *sensemaker* and *forager,* corresponding to the two sub-loops in the Pirolli and Card sensemaking loop (Pirolli and Card 2005; Vogt et al., 2011). All integrated workspace groups developed these roles. Additionally, group J3 adopted the sensemaker and forager roles. As seen in [Table 3], they collaborated the most out of the independent workspace groups. Even though they had established roles, they still broke off to conduct their own separate analysis and communicated less than integrated workspace groups. Additionally, individual territories developed for group J3, with one participant conducting the overwhelming majority (1365 out of 1420) of mouse clicks in the bottom-right window. The other participant was more spread out throughout the other windows of the display. Thus, while user role development appears to be important for close collaboration or communication, territorial style may be even more important for analytic performance.

5.4. Verbal Agreements

Through analyzing utterances by participants, we observed that integrated workspace groups tended to communicate more and use plural pronouns than the independent workspace groups that worked more quietly and used more singular pronouns. Communication is an important part of establishing common ground, "the knowledge that enables [collaborators] to communicate and, more generally, to coordinate their activities" (Chuah, 2003), through the process of "grounding" to ensure that a successful transaction has taken place (Clark, 1991).

Participants used verbal acknowledgments to confirm that information had been received and that they were on the "same page" as each other:

Participant A: *"Do we have him written down too?"* (*Points at a name in a list*)
Participant B: *"We do."*
They continue to check a few more names and double check that they have taken notes on the individuals.

In addition to confirming a joint understanding of pieces of information, different categories of pronouns were used to denote known information. Plural pronouns (e.g. "we," "our") were often used to denote knowledge that was assumed to be common to both participants, although this rule was not strictly followed (emphasis added):

Participant A: *"**We** saw something about 150 thousand, right?"*
Participant B: *"Yeah, **I** saw that."*

The integrated workspace groups tended to communicate more and use plural pronouns than the independent workspace groups that worked more quietly and used more singular pronouns. Greater amounts of direct communication as well as discussion of common knowledge tended to allow participants to maintain an awareness of the other person's thoughts regarding current hypotheses.

In other words, there were more points of common ground between the participants in the integrated workspace groups. The independent workspace groups discussed commonly known information less frequently than the integrated workspace groups. This is also linked to lower levels of close collaboration in the independent workspace groups [Table 3]. Ultimately, the integrated workspace groups had higher, although not significantly so, scores than the independent workspace groups. This difference may be linked to higher levels of communication as well as more closely coupled collaboration.

5.5. Solution commonalities

Territorial styles and levels of communication were also linked with the similarity, or rather, dissimilarity of the individually written reports. Note that these reports were not used in scoring the overall performance of the teams, which was based on the verbal debriefing done after writing the reports. These similarity scores, calculated by summing the common entities and hypotheses reported and subtracting what was only reported by one of the participants, can be found in [Table 3].

Integrated workspace groups tended to have more similar reports than individual workspace groups, indicating that more shared understanding of entity importance and relevance was established between groups that treated the large, high-resolution display as an entirely shared space.

The similarity of individually reported solutions is linked with the correctness of the overall solution, which indicates that groups that came to a joint solution were more accurate with their final hypothesis than those whose solutions diverged from one another.

In summary, increased "sharedness" increases the accuracy of collaborative sensemaking for intelligence analysis using large, high-resolution displays, which is in line with the findings in (Isenberg et al., 2010), where the closeness of collaboration is strongly linked with solution quality. The groups that worked closely together and did not maintain private partitions on the display were more successful in correctly completing the analysis scenario.

6. Design suggestions

Through our analysis, we were able to extract several design suggestions for designing collaborative tools for unstructured text analytics on a large, high-resolution vertical desktop display. These should be treated as suggestions rather than strict guidelines as further studies and analysis may be needed to confirm them.

In order to make use of the available space in a meaningful manner, we suggest that new documents should be displayed in new windows. This allows users to spread information out on the display in order to emphasize the externalization and schematization of information, establish points of shared information, support parallel reading of documents, and re-find information through physical navigation instead of virtual navigation. This feature could be implemented in a variety of ways. For instance, a simple toggle could be used to open new instances of documents, or documents could be dragged and dropped onto a whitespace area to open a new document view.

To support higher levels of communication and shared knowledge, we suggest encouraging an integrated workspace approach to the display, thus emphasizing group territories. However, this suggestion appears to contradict Scott et al.′s work which describes the usefulness of three types of territories: group, storage, and individual (Scott 2004). Individual territories allow persons to conduct analysis and develop hypotheses without premature judgment from other group members. With only group and storage territories, collaborators may run into the problem of groupthink (Janis 1972). Therefore, we do not recommend that individual territories be eliminated. Rather, we suggest that group territories and communication are encouraged.

Multiple independent mice open the possibility of large shared territories on a large, high-resolution vertical desktop display. There was a marked difference in mouse activity between territorial styles (based on where each user′s clicks were distributed). This could be used to automatically detect collaborative or independent activity and adjust affordances accordingly.

7. Conclusion

Having a large, high-resolution display as a workspace allowed the pairs to work in a co-located setting with ample room to sit side-by-side without bumping chairs into each other. This, combined with the high resolution of the display, which allowed many documents to be displayed in their entirety, gave users the flexibility to establish shared, individual, and storage territories in various capacities (Scott 2004). However, we discovered that groups that tended towards an entirely shared display space were more successful in their analysis.

Analytical tools designed to naturally expand into the display space allowed the participants to externalize information to the display in meaningful schemas. The document viewer displayed static document contents in persisted locations on the display, while Jigsaw dynamically updated document contents often in one or two document viewers (one per person). This should not be interpreted to refute the benefits of visualizations for sensemaking. Rather, the affordances of multi-window visualizations may influence space utilization. Methods of organizing information on the display did not appear to directly impact sensemaking performance. Instead, collaboration style and territorial behavior appeared to make a greater impact on performance. Higher levels of "collaborativeness" were linked with higher performance. Effective means of document organization for establishing common ground appeared to promote this collaborativeness.

Participants were able to use the large display to re-find information using spatial references and physical navigation to achieve conceptual common ground, which was enabled by spatial common ground, or common ground regarding the location of information on the display space. This spatial reference was often accompanied by physical pointing to the display. Thus, large, high-resolution displays appear to be useful for co-located collaborative intelligence analysis by providing this physical data representation space.

Through our analysis of an exploratory user study on co-located collaborative sensemaking, we have identified several suggestions for designing collaborative visual analytics tools for use on large, high-resolution vertical displays. In order to establish these, we analyzed how the users interacted with the large display as well as how they collaborated with each other. The design suggestions presented in this paper are not meant to downplay the importance of entity extraction and representation. Visualization-based systems become important when document collections become large, as they are a method of data reduction and foraging. Even though a 32-megapixel large, high-resolution display can easy display the entire contents of 50 short text documents, this display would be unable to do the same for thousands of documents. It is for this reason that we encourage developers to maintain entity representations, but place a larger focus on being able to display multiple document contents that users deem important.

It is our hope that these suggestions will inspire tool designers to build tools specifically designed for collaborative use on large, high-resolution displays. We believe that large displays have the potential to increase the frequency of collaboration in real-world settings due to the innate affordances of the technology.

We wish to pursue this line of research further by conducting an ethnographic study of professional intelligence analysts to further determine user requirements for collaborative sensemaking in this domain. Specifically, we wish to determine how large, high-resolution displays can be used to benefit these analysts, and how visualization-centric and document-centric designs can be better integrated for this purpose. Additionally, we wish to further explore how to better promote shared territorial behavior on large, high-resolution vertical displays for co-located collaboration.

Acknowledgments

This research was supported by National Science Foundation grants NSF-CCF-0937133, NSF-CNS-1059398, and NSF-IIS-1218346, and the VACCINE Center, a Department of Homeland Security′s Center of Excellence in Command, Control and Interoperability. We also wish to thank Dr. John Stasko and his students at Georgia Tech for providing us with the Jigsaw software and study design suggestions.

References

Andrews, C., North, C., 2012. Analyst′s workspace: an embodied sensemaking environment for large, high-resolution displays. IEEE Visual Analytics Science and Technology, 123–131.

Andrews, C., Endert, A., North, C., 2010. Space to think: large high-resolution displays for sensemaking. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. ACM. pp. 55–64.

Ball, R., North, C. Bowman, D.A., 2007. Move to improve: promoting physical navigation to increase user performance with large displays. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. ACM. pp. 191–200.

Beaudouin-Lafon, M., Huot, S., Nancel, M., Mackay, W., Pietriga, E., Primet, R., Wagner, J., Chapuis, O., Pillias, C., Eagan, J.R., Gjerlufsen, T., Klokmose, C., 2012. Multisurface interaction in the WILD room. Computer. vol. no. 45 (4). pp. 48–56.

Birnholtz, J.P., Grossman, T., Mak, C., Balakrishnan, R., 2007. An exploratory study of input configuration and group process in a negotiation task using a large display. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. ACM. pp. 91–100.

Chuah, M.C., 2003. Visualizing common ground. IEEE Information Visualization, 365–372.

Clark, H.H., 1991. Grounding in communication. American Psychological Association, pp. 127–149.

Elrod, S., Bruce, R., Gold, R., Goldberg, D., Halasz, F., Janssen, W., Lee, D., McCall, K., Pedersen, E., Pier, K., Tang, J. Welch, B., 1992. Liveboard: a large interactive display supporting group meetings, presentations, and remote collaboration.

In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. ACM. pp. 599–607.
Endert, A., Fiaux, P., North, C., 2012a. Semantic interaction for sensemaking: inferring analytical reasoning for model steering. In: Proceedings of the IEEE Transactions on Visualization and Computer Graphics. vol. no. 18 (12). pp. 2879–2888.
Endert, A., Bradel, L., Zeitz, J., Andrews, C., North, C., 2012b. Designing large high-resolution display workspaces. In: Proceedings of the International Working Conference on advanced Visual Interfaces. ACM. pp. 58–65.
George Chin, J., Kuchar, O.A., Wolf, K.E., 2009. Exploring the analytical processes of intelligence analysts. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. ACM. pp. 11–20.
Green, T.M., Ribarsky, W., Fisher, B., 2009. Building and applying a human cognition model for visual analytics. In: Proceedings of the IEEE Information Visualization. vol.no. 8 (1). pp. 1–13.
Gutwin, C., 2008. Supporting informal collaboration in shared-workspace groupware. Journal of Universal Computer Science, 14 (9). 1411-1434.
Gutwin, C., Greenberg, S., 1998. Design for individuals, design for groups: tradeoffs between power and workspace awareness. In: Proceedings of the ACM conference on computer supported cooperative work. ACM. pp. 207–216.
Heer, J., 2008. Design considerations for collaborative visual analytics. IEEE Information Visualization, 7 (1). 49-62.
Heuer, R.J., Pherson, R.H., 2010. Structured analytic techniques for intelligence analysis. CQ Press, Washington, DC.
Isenberg, P., Carpendale, S., Bezerianos, A., Henry, N., Fekete, J.D., 2009. CoCoNutTrix: collaborative retrofitting for information visualization. In: Proceedings of the Computer Graphics and Applications. vol. no. 29 (5). pp. 44–57.
Isenberg, P., Tang, A., Carpendale, S., 2008. An exploratory study of visual information analysis. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. ACM. pp. 1217–1226.
Isenberg, P., Fisher, D., Morris, M.R., Inkpen, K., Czerwinski, M., 2010. An exploratory study of co-located visual analytics around a tabletop display. IEEE Visual Analytics Science and Technology, 179–186.
Janis, I.L., 1972. Victims of groupthink: A Psychological Study of Foreign-Policy Decisions and Fiascoes. Houghton Mifflin, Oxford, England.
Kang, Y.-a., Gorg, C., Stasko, J., 2009. Evaluating visual analytics systems for investigative analysis: deriving design principles from a case study. IEEE visual analytics science and technology, 139–146.
Morris, M.R., Lombardo, J., Wigdor, D., 2010. WeSearch: supporting collaborative search and sensemaking on a tabletop display. In: Proceedings of the ACM Conference on Computer Supported Cooperative Work. ACM. pp. 401–410.
Paul, S.A., Morris, M.R., 2009. CoSense: enhancing sensemaking for collaborative web search. Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. ACM. pp. 1771–1780.
Paul, S.A., Reddy, M.C., 2010 Understanding together: sensemaking in collaborative information seeking. Proceedings of the ACM Conference on Computer Supported Cooperative Work. ACM. pp. 321–330.
Pedersen, E.R., McCall, K., Moran, T.P., Halasz, F.G., 1993 Tivoli: an electronic whiteboard for informal workgroup meetings. Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. ACM. pp. 391–398.
Pickens, J., Golovchinsky, G., Shah, C., Qvarfordt, P., Back, M., 2008. Algorithmic mediation for collaborative exploratory search. In: Proceedings of the ACM SIGIR Conference on Research and Development in Information Retrieval. ACM. pp. 315–322.
Pirolli, P., Card, S., 2005. The sensemaking process and leverage points for analyst technology as identified through cognitive task analysis. In: Proceedings of the International Conference on Intelligence Analysis.
Plaisant, C., Grinstein, G., Scholtz, J., Whiting, M., O′Connell, T., Laskowski, S., Chien, L., Tat, A., Wright, W., Gorg, C., Liu, Z., Parekh, N., Singhal, K., Stasko, J., Evaluating visual analytics at the 2007. VAST symposium contest. Computer Graphics and Applications. vol. no. 28 (2). pp. 12–21.
Rekimoto, J., 1998. A multiple device approach for supporting whiteboard-based interactions. Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. ACM. pp. 344–351.
Robinson, A., 2008. Collaborative synthesis of visual analytic results. IEEE Visual Analytics Science and Technology, 67–74.
Scott, S.D., Sheelagh, M., Carpendale, T., Inkpen, K.M., 2004. Territoriality in collaborative tabletop workspaces. In: Proceedings of the ACM Conference on Computer Supported Cooperative Work. ACM. pp. 294–303.
Stasko, J., Görg, C. and Liu, Z., 2008. Jigsaw: supporting investigative analysis through interactive visualization. IEEE Information Visualization. vol. no. 7 (2). pp. 118–132.
Stefik, M., Foster, G., Bobrow, D.G., Kahn, K., Lanning, S., Suchman, L., 1987. Beyond the chalkboard: computer support for collaboration and problem solving in meetings. Communications of the ACM, 30 (1). 32-47.
Stewart, J., Raybourn, E.M., Bederson, B., Druin, A., 1998. When two hands are better than one: enhancing collaboration using single display groupware. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. ACM. , 287-288.
Stewart, J., Bederson, B.B., Druin, A., 1999. Single display groupware: a model for co-present collaboration. In: Proceedings of the SIGCHI conference on human factors in computing systems. ACM. pp. 286–293.
Stewart, J.E., 1997. Single display groupware. In: Proceedings of the SIGCHI conference on human factors in computing systems. ACM. pp. 71–72.
Thomas, J., Cook, K., 2005. Illuminating the path: the research and development agenda for visual analytics. IEEE Computer Society Press, Los Alamitos, (CA,) USA.
Tobiasz, M., Isenberg, P., Carpendale, S., 2009. Lark: coordinating co-located collaboration with information visualization. In: Proceedings of the IEEE Transactions on Visualization and Computer Graphics. vol. no. 15 (6). pp. 1065–1072.
Vogt, K., Bradel, L., Andrews, C., North, C., Endert, A., Hutchings, D., 2011. Co-located collaborative sensemaking on a large high-resolution display with multiple input devices. In: Proceedings of the Conference on human-computer interaction. Springer. pp. 589–604.