

Designing Display Ecologies for Visual Analysis

Haeyong Chung

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

Doctor of Philosophy
in
Computer Science and Applications

Chris North, Chair
Doug Bowman
Niklas Elmqvist
Steve Harrison
Ben Knapp

Feb 20, 2015
Blacksburg, Virginia

Keywords: display ecology, multi-display environment, ubiquitous analysis, interactive visualization, visual analytics, sensemaking

© Copyright 2015, Haeyong Chung

Designing Display Ecologies for Visual Analysis

Haeyong Chung

ABSTRACT

The current proliferation of connected displays and mobile devices—from smart phones and tablets to wall-sized displays—presents a number of exciting opportunities for information visualization and visual analytics. When a user employs heterogeneous displays collaboratively to achieve a goal, they form what is known as a display ecology. The display ecology enables multiple displays to function in concert within a broader technological environment to accomplish tasks and goals. However, since information and tasks are scattered and disconnected among separate displays, one of the inherent challenges associated with visual analysis in display ecologies is enabling users to seamlessly coordinate and subsequently connect and integrate information across displays. This research primarily addresses these challenges through the creation of interaction and visualization techniques and systems for display ecologies in order to support sensemaking with visual analysis.

This dissertation explores essential visual analysis activities and design considerations for visual analysis in order to inform the new design of display ecologies for visual analysis. Based on identified design considerations, we then designed and developed two visual analysis systems. First, VisPorter supports intuitive gesture interactions for sharing and integrating information in a display ecology. Second, the Spatially Aware Visual Links (SAViL) presents a cross-display visual link technique capable of guiding the user's attention to relevant information across displays. It also enables the user to visually connect related information over displays in order to facilitate synthesizing information scattered over separate displays and devices. The various aspects associated with the techniques described herein help users to transform and empower the multiple displays in a display ecology for enhanced visual analysis and sensemaking.

DEDICATION

For my mom and dad,

BokSoon Kim (김복순) and HwaJin Chung (정화진 1935-2011)

Table of Contents

ABSTRACT	ii
DEDICATION	iii
Table of Contents	iv
List of Tables	viii
List of Figures	ix
1 Introduction.....	1
1.1 Research Overview.....	2
1.2 Thesis Contributions	4
1.3 Dissertation Outline.....	5
2 Background and Related Work.....	7
2.1 Display Ecologies.....	7
2.1.1 <i>The concept of ecology in the context of HCI research.....</i>	<i>7</i>
2.1.2 <i>Display Ecologies for Visual Analysis</i>	<i>10</i>
2.2 Visual Analysis Systems and Techniques.....	11
2.2.1 <i>The Value of Space for Sensemaking.....</i>	<i>11</i>
2.2.2 <i>Visual Analysis Tools for Single Users</i>	<i>12</i>
2.2.3 <i>Collaborative Visual Analysis Tools.....</i>	<i>13</i>
2.2.4 <i>Visual Analysis tools on Emerging Displays.....</i>	<i>14</i>
2.2.5 <i>Cross-display Interaction and Visualization Techniques.....</i>	<i>15</i>

2.2.6	<i>Multi-Display Systems and Environments</i>	18
2.2.7	<i>Software Frameworks for Multi-display Environments</i>	19
3	Design Considerations for Visual Analysis in Display Ecologies	21
3.1	Method.....	22
3.2	A Scenario for Visual Analysis in Display Ecologies.....	24
3.3	Design Considerations for Display Ecologies.....	25
3.3.1	<i>Display Composition</i>	27
3.3.2	<i>Information Coordination</i>	30
3.3.3	<i>Information Connection</i>	33
3.3.4	<i>Display Membership</i>	37
3.5	Discussion	40
3.5.1	<i>Balance Foraging and Synthesis Approaches</i>	41
3.5.2	<i>Exploit the Physicality of Displays</i>	42
3.5.3	<i>Provide Spatial Inter-Awareness of Displays</i>	44
4	VisPorter: Facilitating Information Sharing for Collaborative Sensemaking in Displays Ecologies	45
4.1	Design Goal	47
4.2	The VisPorter System Overview.....	49
4.2.1	<i>Usage Scenario</i>	50
4.2.2	<i>Sensemaking Tools</i>	55
4.2.3	<i>Display Proxy Interface</i>	57
4.2.4	<i>Gesture-based Interaction</i>	58
4.2.5	<i>Implementation</i>	60
4.3	Evaluation	61
4.3.1	<i>Participants</i>	61
4.3.2	<i>Task</i>	62
4.3.3	<i>Apparatus</i>	62
4.3.4	<i>Procedures</i>	63
4.3.5	<i>Data Collection and Analysis</i>	63

4.4	Findings.....	65
4.4.1	<i>Collaboration Styles with Multiple Displays</i>	65
4.4.2	<i>Cross-Display Semantic Structures</i>	67
4.4.3	<i>On-demand Extension of Display Space</i>	69
4.4.4	<i>Objectification of Information</i>	71
4.4.5	<i>User Feedback</i>	72
4.5	Discussion.....	74
4.5.1	<i>Performance Factors</i>	74
4.5.2	<i>Deciding Better Analysis Strategies</i>	75
4.5.3	<i>Spatial and Physical Actions</i>	77
4.5.4	<i>Opportunistic Activities</i>	78
4.5.5	<i>Promoting the Objectification of Information</i>	79
4.6	Summary.....	79
5	A Comparison of Two Display Models for Collaborative Sensemaking	81
5.1	Two Display Models.....	83
5.1.1	<i>VizCept: Shared Visualization Spaces</i>	84
5.1.2	<i>VisPorter: Display Ecology</i>	85
5.2	Study Description.....	87
5.3	Findings.....	89
5.4	Discussion.....	94
5.5	Summary.....	96
6	SAViL: Spatially-Aware Visual Links for Sensemaking in Display Ecologies	97
6.1	The SAVIL Overview.....	100
6.1.1	<i>Cross-Display Visual Links</i>	102
6.1.2	<i>The SAViL Drawing Algorithm</i>	106
6.1.3	<i>Prototype System and Implementation</i>	108
6.1.4	<i>Usage Scenarios</i>	111
6.2	User Study.....	113
6.2.1	<i>Participants</i>	114

6.2.2	<i>Dataset and Task</i>	114
6.2.3	<i>Apparatus</i>	115
6.2.4	<i>Procedures</i>	115
6.2.5	<i>Data Collection and Analysis</i>	116
6.3	Observations.....	116
6.3.1	<i>Visual Link Usage Observations</i>	116
6.3.2	<i>Information Foraging and Awareness</i>	117
6.3.3	<i>Help Leveraging Multiple Displays</i>	118
6.3.4	<i>Semantic Structures across Displays</i>	119
6.3.5	<i>Synthesis across Displays</i>	120
6.4	Discussion	124
6.5	Summary	126
7	Conclusion	127
7.1	Restatement of Contributions.....	127
7.2	Limitations and Future Work.....	130
7.2.1	<i>The Studies</i>	130
7.2.2	<i>Automatic View Adaptation for Multiple Displays</i>	133
7.2.3	<i>Support for Software Framework and Infrastructure</i>	134
7.2.4	<i>Analysis Provenance for Display Ecologies</i>	135
	Final Remarks	135
	References	136

List of Tables

Table 3.1. Design considerations of display ecologies for visual analysis.	26
Table 3.2. Display memberships.....	37
Table 4.1. Study result.	64
Table 5.1. Design characteristics of the two systems.	86
Table 5.2. Key differences between the two models.....	95
Table 6.1. Evaluation results.....	114

List of Figures

Figure 2.1. Space to think: large displays for sensemaking [1]	12
Figure 3.1. Four key design considerations for visual analytics in display ecologies.	22
Figure 3.2. Tangible views and a vocabulary of physical movement of cardboard lenses [52].	29
Figure 3.3. Examples for semantic substrates in display ecologies. left: affinitytable [84]. right: design studio [83].	30
Figure 3.4. Seamless cross device interaction to move objects from one device to another – Throwing interface [87].	33
Figure 3.5. The pixel-oriented treemap [95]......	39
Figure 3.6. Exploit affordance of interaction for multiple displays. The stackable interface (left) [2] and moveable focus+context displays (right) [3].	43
Figure 4.1. VisPorter is a collaborative text analytics tool for multiple displays.	47
Figure 4.2. Foraging tool - Document viewer.....	52
Figure 4.3. Foraging tool - ConceptMap viewer.....	52
Figure 4.4. Two types of document boxes for the synthesis tool (a) text document and (b) image.....	53
Figure 4.5. Synthesis tool on the shared display.....	54
Figure 4.6. Easy to connect between two entity nodes by tapping gestures.....	56
Figure 4.7. Swipe and drop the document onto the shared displays: (a) Wall displays and (b) Tabletop display.....	59

Figure 4.8. Transfer and merge individual concept maps and entities in a wall display through tap-holding gestures.	59
Figure 4.9. Three collaboration styles for multiple displays. Blue arrows indicate users.	67
Figure 4.10. Organizing information based on multiple entity types on different displays. On the figure of the wall display, we added labels pertaining to participant explained regions of clustered documents described to us during the debriefing.....	69
Figure 4.11. User feedback in the post-session survey (1-5 scale).....	73
Figure 4.12. Cross-device referencing with physical navigation. The user in G4 analyzed the concept map on his iPad and text documents on the tabletop. He used physical navigation to scan the documents on the tabletop rather than use the search feature.	78
Figure 5.1. The two collaborative sensemaking systems used in our comparative studies.	82
Figure 5.2. Two display models for collaborative sensemaking.	85
Figure 5.3. The common analysis workflow.....	90
Figure 6.1. Spatially aware visual links for display ecologies.	99
Figure 6.2. SAViL with basic document analysis tools: (a) word cloud, (b) document search interface, (c) highlighting and shoeboxing interface, and (d) document artifact.....	101
Figure 6.3. SAViL cross-display links. Each rounded box represents a document, and small red and yellow boxes represent entities. A user clicks an entity in a document on Display B and every same entity on different displays is automatically connected.	102
Figure 6.4. Manual linking. From left-to-right: (a) a user drags the anchor across two displays, (b) place it on a target document, and (c) a manual link is drawn across the displays.....	104
Figure 6.5. Support spatially aware links. A small display around a tabletop display is moved to a different location and the cross-display links keep following the location of the display.....	105

Figure 6.6. Drawing SAViL from 3D physical space to 2D screen. Green boxes represent the displays laid in the space, the red dots represent the position of each display in the space, blue dotted lines represent the 3d virtual lines, and solid red lines represent projected visual links on each display. 106

Figure 6.7. SAViL client/server architecture..... 110

Figure 6.8. U6 organized documents based on chronological order, effectively building a timeline of events with the earliest events. 122

Figure 6.9. Three different plots across the displays. We added labels pertaining to participant explained different subplots and clustered documents across displays, which were described to us during the debriefing. The different colored regions represent different..... 123

Figure 7.1. An example of a line chart automatically rendered at different scales. The algorithm preserves the various elements of the line chart based on their semantic importance at a given scale. 134

1 Introduction

Rapidly advancing technologies signify our increasing access to various types of devices for personal and professional use—from smart phones, tablets, and laptops to desktop PCs and large high-definition displays. In fact, these various displays are helping us accomplish ever more challenging tasks and goals, which seem to be propagating with the availability of larger data sets that require novel analytical approaches in order to manage and interpret them. While these multiple displays may be collocated in a workspace or living area, in most cases they function as independent display screens. However, when a user employs heterogeneous displays collaboratively to achieve a goal, they form what is known as a *display ecology*. The display ecology enables multiple displays to function in concert within a broader technological environment to accomplish tasks and goals. Among various domains potentially supported by display ecologies, this research focuses on visual analysis applications for supporting and externalizing the process of “sensemaking” [4].

Display ecologies can better assist people in enhancing visual analysis with larger and discretized display space for analysis, which is augmented by various interaction affordances facilitated by the different displays [5], [6]. However, since information and tasks are scattered and disconnected among separate displays, one of the inherent design challenges associated with visual analysis in a display ecology is enabling users to seamlessly *coordinate* and subsequently *connect* and *synthesize* information across displays. This work primarily addresses these essential challenges through the design of new visual analysis techniques and systems for display ecologies in order to support the analysis of large, complex document datasets.

1.1 Research Overview

The display ecology provides a new design opportunity that must consider how multiple displays and devices can be used for analysis of the constantly escalating quantity of data. How can we design more efficient and supportive visual analysis tools with multiple displays and devices? A central goal of this research is to investigate the problem of how to design and develop the visualization tools and techniques enabling multiple displays to function within display ecologies in order to realize their full power in support of visual analysis. In particular, we investigate the cross-device interaction and visualization techniques that can best leverage the discretized display space and interaction affordances facilitated by different displays for sensemaking tasks. The design of our tools and techniques are primarily grounded in two theories: distributed cognition [7] and the concept of *Space to Think* [1]. Thus, our visual analysis systems are designed for users to perform data analysis by embedding information and analysis components across different displays. Specifically, the ecology systems focus on how multiple displays enable space to think where users employ the discretized screen space as (1) semantic structure and (2) external memory.

This research contributes to visual analytics tools based on display ecologies by addressing three main research questions and their sub-questions:

- 1) How can a visual analysis tool be designed to leverage a display ecology?

- a. What are the core analysis activities that inform the design of visual analysis tools in a display ecology?
 - b. What are the design considerations of display ecologies for visual analysis?
- 2) What interaction and visualization techniques are needed in order for a display ecology to offer the same "space-to-think" benefits [1] as large displays with readily accessible devices?
- a. How can users coordinate and organize analysis tasks and information among multiple displays?
 - b. What visualization techniques can help users connect and synthesize scattered information among data items spread across displays?
- 3) What are the effects of the presented techniques and systems on sensemaking in a display ecology?
- a. How does a display ecology impact a user's sensemaking and thought processes while analyzing a large number of text documents?
 - b. How can users externalize their sensemaking process to a display ecology?

Our exploration of the first question is concerned with identifying and creating salient dimensions of design considerations and associated visual analysis techniques and tools for a display ecology. Identifying design principles to support visual analysis is complex because a display ecology presents a new set of challenges, as well as a paradigm that was not considered during the design of visual analytics tools intended for use on single displays.

In order to answer the second research question, we investigate novel interaction and visualization techniques intended to provide a cohesive and integrated analysis experience for users in a display ecology—emphasizing seamless analysis experience among heterogeneous displays. The techniques and systems focus on visual text analytics that enable users to distribute information and analysis components onto various displays.

To address the third research question, we conduct user studies for evaluating the presented techniques and systems for visual analysis in display ecologies. In these studies, we investigate how our presented techniques and display ecologies can impact the strategy and process of sensemaking.

The primary focus of this research entails designing, implementing, and evaluating novel visualization techniques and systems that utilize multiple displays and their interaction affordances enhancing a user's ability to refine and comprehend important information hidden in large, complex datasets.

1.2 Thesis Contributions

This research contributes new systems and techniques that enable sensemaking with multiple displays. In particular, this research will increase our understanding of the ways in which seamless interaction experiences can heighten the effectiveness of integrating multiple displays as a single ecology for visual analysis. Such knowledge will guide the creation of more appropriate visual analysis tools for multiple displays and ubiquitous environments. The anticipated contributions can be categorized in three areas:

- 1) Refining the design considerations for visual analysis in a display ecology.
 - a. Identification of core visual analysis activities in a display ecology.
 - b. Exploration of the comprehensive design choices for visual analysis tools in a display ecology.
- 2) Tools and techniques to support visual analysis and sensemaking in display ecologies.
 - a. Creation of display ecology systems directed toward supporting sensemaking, emphasizing external memory and semantic structures.
 - b. Development of interaction and visualization techniques that allow multiple users to physically *share* and *synthesize* both information and visualizations across displays.
- 3) Evaluating the presented tools and techniques for the visual analysis process.

- a. Understanding the diverse impacts of a display ecology on visual analysis and the sensemaking process.
- b. Investigating the effectiveness of suggested features for sensemaking in a display ecology in terms of semantic structures and external memory.

1.3 Dissertation Outline

Chapter 2 describes the overview of research in the concept of ecologies and visual analysis research in the model of sensemaking process. This includes defining display ecology. We also survey and categorize the state of the art in multi-display systems.

Chapter 3 introduces design considerations for visual analysis tools based on display ecologies. This chapter describes the core analysis activities that was employed to analyze the design consideration of a display ecology for visual analysis and explores design considerations based on the analysis activities. The design of the visual analysis tools presented in this dissertation are guided by these design considerations.

Chapter 4 introduces *VisPorter*, which supports intuitive gesture interactions for sharing and integrating information in a display ecology. Essentially, *VisPorter* enhances analysis tasks by enabling users to distribute information across multiple personal devices (e.g., smart phones, tablets, etc.) and shared displays (e.g., wall displays, etc.). *VisPorter* emphasizes the importance of immediacy in information sharing and synthesis across displays by implementing a gesture-based interface.

Chapter 5 describes a comparison of the use of two display models for visual analysis, one based on the model of the personal displays with shared visualization spaces (*VizCept* [8]) and the other based on the distributed model whereby different displays can be appropriated as workspaces in a unified manner by collocated teams (*VisPorter*).

Chapter 6 introduces *Spatially Aware Visual Links (SAViL)*, a cross-display visual link technique capable of (1) guiding the user's attention to relevant information across displays; and (2) visually connecting related information among displays. *SAViL* visually

represents the connections between different types of information (e.g., keywords, documents, pictures, etc.) across displays.

Chapter 7 concludes this dissertation by summarizing the work and contributions and providing suggestions for future research.

2 Background and Related Work

2.1 Display Ecologies

In this section, we define and describe the key ecology concepts informing our design of display ecologies and discuss how display ecologies can support visual analysis.

2.1.1 The concept of ecology in the context of HCI research

The Oxford English Dictionary defines “Ecology” as “*The branch of biology that deals with the relations of organisms to one another and to their physical surroundings.*” Biologists use the term ecology to describe interconnections within our natural world – a community of living organisms in conjunction with the nonliving components of their environments (e.g., air, water, minerals, soil, etc.), interacting as a system. As research expands across the disciplines, the notion of ecology has migrated to other areas beyond its biological definition. In a variety of fields, including the social sciences and HCI, the concept of ecology has been broadly employed to describe and understand the spaces or

environments in which individuals or groups of people work or interact with technologies, guided by their own goals and values.

Different ecologies emphasize different aspects of modern technological environments. For example, Krippendorff investigated the ecological meaning of a community of artifacts [9]. Specifically, he asserted that a collection of artifacts forms an ecology, which inevitably creates different relationships with other artifacts (of other types) such as cooperation, competition, and independence; these then continuously guide particular users' choices, driving an increase or decrease of species (types) of artifact. Crabtree and Rodden described a hybrid ecology, which combines mixed reality environments and virtual environments [10]. They regarded this type of ecology as the space or specific environment in which physical and digital media are socially organized and users can take advantage of them to achieve individual activities. Nardi and O'Day [11] emphasized an ecology as a system of diverse technologies, people, and practices whereby information as nutritive elements or energies is produced, consumed, accessed, transited and circulated among the different technologies and devices based upon users' different tasks. Despite their nuanced differences, these ecology concepts share the same foundation in that they facilitate an enhanced understanding of complex dynamics, relationships and interactions among users, technologies, and work practices—all fundamentally based on some ecological processes.

We extend these various notions of ecologies by applying them to the modern technological concept of a multi-display environment. When users employ available displays collaboratively to achieve a goal, they combine available displays to accomplish their desired outcomes. A group of such displays (e.g., a smart phone + computer + HDTV + large screen display) form an ecology in which the displays can relate to one another in a variety of ways. The basic premise, however, is that each display plays a different role in the workflow for specific goals, consuming information as nutritive elements. Therefore, A *Display Ecology* is defined as a system of heterogeneous displays that engage the entire workflow of a task to better assist people in achieving their desired outcomes.

The key aspect of incorporating the “ecology metaphor” in designing applications for multiple displays is to target the relationships and interaction among various displays—rather than designing individual visual and interaction techniques for each display. Such an approach provides essential design considerations based on natural ecosystems about how information, practices, artifacts, and heterogeneous technologies should coalesced in a holistic system. Thus, designing a display ecology is germane to determining how multiple displays mutually interact, support, and collaborate with one another to solve user’s specific problems.

There are several prior studies related to display ecologies in the domain of HCI research. They are largely concerned with how an ensemble of displays can better assist people in fulfilling various applications given the differing characteristics of the system, such as larger and discretized display space and various interaction affordances facilitated by different displays. Huang et al. conducted a field evaluation of the large display ecology used in the NASA Mars Rover exploration mission [5]. They described how the role of displays and the collaboration style of an ecology can change over time. Their results also strongly support the positive opportunistic characteristics of such an ecology. Coughlan et al. [12] conducted a study examining university students’ fieldwork using multiple laptops, tabletops and projectors. Based on their results, they presented an ecology framework for analyzing relationships between displays, such as seams, bridges, niches, and focal characters. Terrenghi et al. provided a taxonomy for the scale of the ecology, the nature of social interaction, and the interaction techniques employed [13]. These three aspects of a display ecology also confirm spatial, semi-collaborative, and opportunistic characteristics.

A few research projects have informed advanced design considerations for multiple displays. Waldner et al. [14] provided design considerations for a multi-view visualization based on a multiple display environment; specifically, they described appropriate interaction techniques to facilitate information sharing, the design of flexible working environments, and so forth. Badam et al. [15] formally define several design

considerations for multi-display visualizations based on the composite visualization framework [16].

2.1.2 Display Ecologies for Visual Analysis

In an information-intensive world, new approaches to visual analysis are needed to enable people to benefit from the availability of massive, complex information. This dissertation focuses on the design and development of new visual analysis tools, informed by the ecological implications of a community of available displays. The versatility of an ecology of devices and displays has potential for significantly impacting visual analysis that emphasize the formation of insight through larger screen space and various interaction affordances facilitated by different displays. We can consider the three main characteristics of a display ecology for visual analysis.

First, display ecologies provide larger display spaces beyond one single virtual raster space and they also enable users to increasingly utilize multiple screen spaces as a resource for visual perception and spatial ability. Display ecologies also facilitate better utilization of physical space because separate displays can be located at different angles, as well as in different places. The larger physical space afforded by display ecologies can play an important role in insight formation. Indeed, prior research has shown that the use of large physical spaces impacts insight formation significantly. For example, a large screen real estate and greater resolution enable a user to visualize larger amounts of data, as well as provide more space for *physical navigation* choices that effectively exploits human spatial senses and embodied cognition. Ball et al. [17] confirmed that physical navigation produced a performance improvement in visualization tasks over virtual navigation.

Second, multiple displays are able to satisfy the analytic needs of both individual and multiple users in a group. A display ecology's semantically discretized space for analysis enables multiple users to divide their data and tasks among different displays. Multiple discretized screen spaces of personal displays facilitate the division of analysis task into small parts, thereby enabling users to focus on visualizing one major piece of the dataset in one display, or its related information across the displays. In particular, this

characteristic of display ecologies allows collaborating users to distribute individual and collaborative work effectively with and among individual and shared displays. Visual analysis is often the combined result of individual and collaborative efforts. Even though users may be working on a collaborative analysis, individual tasks remain important; in fact, users typically spend more time on individual analyses, even during collaborative sessions [18]. With display ecologies, individual users can solve a given problem independently on each individual display while seamlessly collaborating with others on the shared displays.

Lastly, Display ecologies can accommodate an analyst's changing needs by enabling the user to combine or shift different displays and devices, tapping into the potential of different types of technologies for suitable tasks [12], [5], [19]. During data analysis, analysts encounter, carry, and consult various pieces of information at opportunistic moments. They may emerge from the domain knowledge of the analyst by chance [20]. The ecology approach may enable continuous capture of the insight formation process as it occurs in different contexts.

2.2 Visual Analysis Systems and Techniques

This research also drew inspiration of visual analysis systems for display ecologies from visual text analytics tools for both small and large displays, multi-display environments, various cross-display interaction and visualization techniques. In this section, we review existing visual analysis projects and study results.

2.2.1 The Value of Space for Sensemaking

We extend prior studies examining the value of space for sensemaking. Robinson et al. [21] found that analysts conducting sensemaking tasks on a set of short text documents printed on notecards used table space to spatially organize the documents. Andrews et al. [1] found that large display systems enabled a similar phenomenon that they called “space to think”, in which users utilized the additional screen space to support semantic structure and external memory. Essentially, users used the space to spatially organize and structure the information and their thoughts. Their studies emphasize spatial organization of

various documents and entities, enabling the analyst to leverage the larger screen space for rapid externalization of their cognitive syntheses during the sensemaking process (Figure 2.1). They found that strictly virtual space on small physical screens did not consistently produce such behaviors, indicating that the large physical space was critical to enabling the behavior [22], [23]. Hamilton et al. [24] also found similar results when users employed a large number of small mobile displays. Specifically, they observed that users organized the physical devices on a table in order to spatially structure information displayed on those devices.

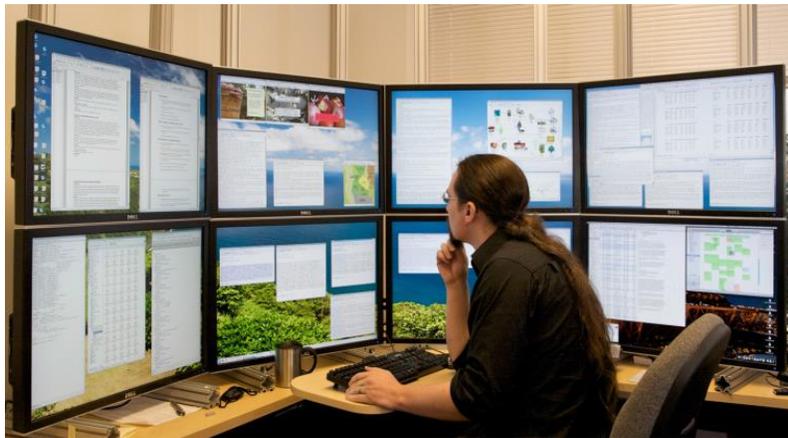


Figure 2.1. Space to think: large displays for sensemaking [1].

2.2.2 Visual Analysis Tools for Single Users

The presented visual analysis systems in this research combined features of visual analytics for single displays in order to support various text analytic activities in display ecologies. *Sandbox* in the nSpace suite is designed to support an open workspace where users can move information objects and organize on the display space for external representations [25]. Andrews et al. expanded the benefits of the external representation to sensemaking tasks on personal large displays [23]. Our research was also motivated by *Jigsaw* [26] in that it provides visual illustrative connections between automatically extracted entities in multiple documents. Analyst's Notebook by i2 Inc. provides semantic graph visualization for link analysis to identify connections and patterns in a large

aggregate dataset [27]. It allows users to visualize and analyze large quantities of intelligence data through basic link and node charts.

Our proposed systems (described in Chapters 4 and 6) were motivated by these systems as a way to spatially organize evidence and resulting hypotheses across multiple screens.

2.2.3 Collaborative Visual Analysis Tools

Another area related to this work is that of collaborative visual analysis systems. Many of these collaboration tools focus on group exploration of data through community components such as annotations and comments. They emphasize the engagement of users in data analysis through various types of social navigation such as associated discussion components where users post comments or annotations, and ask questions.

Entity Workspace is modelled on the effectiveness of a traditional evidence file that keeps track of various facts about entities and relationships, such as people, places, organizations, telephone numbers, bank accounts, etc. [28]. It provides an explicit model of important entities by allowing users to find potentially important documents and entities. This tool helps analysts rapidly find the new facts based on these connections of entities and information. In their follow-up work on Entity Workspace, Bier et al. developed five design guidelines for collaboration for intelligence analysis and modified the Entity Workspace system based on these guidelines [29]. The modified tool helps collaborators to connect entities and concepts that are found by different analysts, allowing them to be merged and shared seamlessly. Pike et al. developed a service oriented visual analytics system, SRS [30], which distributes analysis tasks to allow client applications running on different devices, such as mobile devices and laptops. The SRS web client incorporates web services including manually created concept maps, timeline visualizations, and listings of query results. It also allows users to save and share questions, hypotheses, evidence, etc.

In many collaborative visualizations, one user initially creates the visualization, and other users add annotations or show interesting views for the visualization. For example, *ManyEyes* allows users to upload their data onto a public website and build, share, and

edit information visualizations from the data [31]. Another example is *Dashiki*, a wiki-based website which enables multiple users to build wiki-based visualization dashboards through a user-editable wiki mark-up language and interactive editors [32]. Users can present and organize their own dashboards, which contain visualization and presentations that are created by multiple users as community components. Heer et al. developed a web-based asynchronous collaboration visualization tool, *sense.us* [33]. It provides a set of interactive visualization features along with collaboration via bookmarking of views, a new discussion scheme called "doubly-linked", graphical annotation, and social interaction through annotations, comment listings and user profiles. Increasingly, many web-based services are supporting more collaborative visualization features for general web users. For example, Tableau Public supports a personalized visualization service that allows multiple users to share specific visualization and data with others [34] on the web.

2.2.4 Visual Analysis tools on Emerging Displays

More recently, new visual analysis systems based on non-traditional shared screen spaces have begun to support co-located collaboration for visual analysis. For example, *Cambiera* enables multiple users to search and manage documents through its unique widgets and allows them to organize documents collaboratively on the tabletop [35]. Tobias et al. developed a system called '*Lark*' [36] that lets multiple users analyze the same data with visualizations on a tabletop. *Hugin* focuses on enabling multiple remote users to synchronously interact with shared visualizations on large displays [37]. The Branch-Explore-Merge [38] approach allows multiple users to privately modify information on individual displays and then merges their changes onto a shared display upon the agreement of other group members. Jetter et al. [39] presented *Facet-Streams*, a collaborative search tool that allows users to combine multiple search features with a tangible user interface in order to filter a dataset. It utilizes multi-touch interaction and tangible tokens placed on a tabletop to display multiple filter streams that can then be combined into single streams to view the filtered data.

Results of some user studies for visual analysis on these emerging display spaces also inform design implications for this research. Vogt et al. [18] and Isenberg et al. [40]

adapted existing visual analytics tools for multiuser interaction on large display environments. Isenberg et al. [41] addressed the types of collaboration styles that are adopted during co-located collaborative visual analysis on a single tabletop. These studies highlight the importance of integrating visual analysis tools into the document spaces in which users exploit a large display space for collaborative visual analysis.

2.2.5 Cross-display Interaction and Visualization Techniques

The following multi-display systems focus more on information sharing across multiple displays. For example, Wigdor et al. designed an interaction space that uses multiple wall mounted displays and one tabletop display [42]. Each wall mounted display can receive digital objects through the *World in Miniature* (WIM) on the tabletop display, in which a digital object can be placed on a wall mounted display by dragging it to the corresponding WIM. Johanson et al. developed a framework called *multibrowsing* [43] to deploy the plug-in application in *iRoom* that is an interactive space that consists of multiple displays [44]. This system allows users to move Web content among heterogeneous displays including personal laptops and wall-sized displays, but each device plays the role of a full-screen web browser only and doesn't allow users to organize the information spatially across the these displays in order to leverage the large screen spaces.

A few multi-display systems focus on screen sharing to enable users to transfer their private laptop windows onto larger shared displays. Also, input redirection enables users to interact with the shared windows by using any of the devices (private or shared) for input. For instance, *WinCut* allows the user to specify and organize regions of information (ROI) in multiple windows from different laptops [45]. Users can replicate arbitrary regions of selected windows in other separate windows called WinCuts. Those WinCuts allow for visual updates and direct interaction, and they can be shared across different displays. Users can share different WinCuts from different laptops on wall screens. *WeSpace* provides an image sharing system which enables users to share the entire desktop with multiple windows from different laptops and a tabletop in flexible screen layouts [46]. *LivOlay* is a system developed to overlap multiple application windows for data comparison [47]. Remote laptops running different applications are connected to a

large display which is then used to overlay the information. The system employs user-chosen common points to register the images in order for the images to be correctly associated with one another. Interaction with the large display is through a tablet and/or any remote laptop's cursor. LivOlay is beneficial when comparing images, graphs, etc. because it does not alter the original display, although modifications made on the large display are available on the remote laptops. The *Dynamo* system is a public display system designed to support simultaneous multi-user interactions with one or more large displays in public meeting settings [48]. It allows users to transfer media files from their personal devices or laptops to a large display and allows users to exploit these large displays as extensions of their personal devices to enable sharing and exchange of personal media files with other users through the large displays. Likewise, Greenberg et al. presented the *Shared note* system which allows people to move private notes created through PDAs selectively onto public displays [49].

Also, these new systems support handheld devices as a special interface with larger displays or a single primary display for visual analysis or data exploration. *Smarties* [50] is a new input system to facilitate development of wall display applications. The system allows for the prototyping of interactive applications, whereby multiple mobile devices serve as remote controllers interacting with a wall display. Specifically, users can interact with the main wall display application by manipulating multiple interactive *pucks* on their mobile displays, which represent different GUI widgets (e.g., cursors, buttons, sliders, menu, etc.) and content (e.g., visual objects and text) of the wall display. Jansen et al.'s Tangible Remote Controller is also based on both a tablet and tangible user interfaces to interact with wall-sized displays [51].

Several multi-display systems enable users to customize visualization or media views by physically moving or aggregating portable displays in the physical space. Interconnected and spatially aware displays can couple visualization environments and physical environments directly, allowing users to engage and employ physical skills. For example, Spindler et al. [52] presented Tangible views which are cardboard interfaces created using overhead projections in conjunction with a tracking system. This system allows users to

take advantage of physical space and skills by directly moving the cardboard lenses to interact with a large visualization on the tabletop (e.g., Focus+Context, magnification of a piece of the whole visualization, etc.). So movement of the device can be treated as part of the visualization. *i-Loupe* and *iPodLoupe* are techniques developed for use with a tabletop display [53]. Both employ focus + context to visualize and interact with data more efficiently. *i-Loupe* contains two lenses, a base, which allows the user to keep the context within the whole, and a focus, which allows for magnification of a piece of the whole. Using the *iPodLoupe* technique, users can utilize the iPod as the physical focus lens. More recently, Rädle et al. [54] presented *HuddleLamp*, a desk lamp that facilitates spatially-aware interactions around a table by detecting and tracking the movement and position of mobile displays, as well as the actual hands of users with sub-centimeter precision. When users need to reorganize their workspace by adding or removing devices, this system allows for ad-hoc multi-device collaborations and interactions around a table, whereby users can mix and match different available devices.

There are also several systems for creating large, ad-hoc tiled displays from multiple displays to show contiguous views [55], [56]. In these systems, relatively simple media contents, such as an image or a webpage, can be displayed over more than two displays, with navigation possible from either display. These systems consider affordances of different types of displays (mobile and large displays) to create single visualization views.

There are interaction techniques for transferring information across different devices for sharing with other collaborating users [57], [58]. Nacenta et al. provided a taxonomy for the cross-display movement and proposed several categories of interaction techniques for cross-device interaction [59]. Pick-and-drop [60] uses uniquely identifiable interaction devices, such as pens, to transfer digital objects between multiple displays. Dachsel et al. [61] and Marquardt et al. [62] explore cross-device interaction techniques which enable users to tilt devices toward one another for sending information. In this dissertation, we will also present a natural way to transfer and coordinate information and visualization items across different displays and devices (Chapter 4).

Increasingly, the various devices in a display ecology are likely to be co-located. Although they are situated in the same physical space, they may be placed at different angles. Thus, it may be difficult for some users to view all the information with the same clarity, perspective, etc. In response, several perspective window systems have been proposed to provide enhanced visibility and interactivity with the user—regardless of the locations and angles of multiple displays. For instance, *E-conic* [63] was developed to support the dynamic perspective correction of display content and GUI widgets on displays in different locations and at different angles. In this system, windows, information and GUI on multiple displays can be automatically adapted to the differing perspectives of users. *Deskothèque* [64] provides automatic geometric compensation of multiple projector-based displays, which detect and compensate distorted and overlapped displays in order to create a more seamless and planar screen output. Xiao et al. presented the Ubiquitous Cursor [65], which is a system that provides users with direct feedback with respect to the cursor’s location among displays. This system helps users keep track of the cursor across displays with direct visual feedback. To accomplish this, the system employs a projector and a hemispherical mirror where the displays are located.

To evaluate the effectiveness of the system, they conducted a user study that compared the Ubiquitous Cursor with two different feedback approaches for multiple displays: *Halo* [66] and *Stitching* [67]. The results showed that Ubiquitous Cursor was much faster than either of the other approaches for repeated aiming tasks. The Perspective Cursor [68] allows users to map usual mouse cursor to different displays in a multi-display environment based on the user’s perspective view. The technique enables users to be aware of the cursor, which can seamlessly travel across different screens as if the screens were a single desktop PC setting connected to one PC. The system calculates the cursor’s relative position based on the spatial relationships between the user’s head and the location/orientation of each display.

2.2.6 Multi-Display Systems and Environments

There are a few multi-display interactive workspace that share closer design and goals to our presented methods. In display ecology, multiple displays (mostly large stationary

displays) can be located at different places in a room, and an aggregation of their screen spaces increases the screen real estate thus enabling multiple users to see more information. Multiple types of displays are frequently combined to construct multiple coordinated views in workroom or laboratory settings. For instance, Streitz et al. [69] presents *i-Land* where users can mix-and-match multiple portable and large displays and devices. With *i-Land*, users can associate digital objects with physical tangible objects, which can then be used to physically move the objects between computers. Zoomable Object-oriented Information Landscape (ZOIL) [70] provides a framework for designing sensemaking space. ZOIL allows users to freely coordinate documents or visualizations around multiple displays, emphasizing the concept of persistence and external representations. Using a similar concept, Geyer et al. proposed a multi-display system which enables users to organize individual sketches created on between individual tablets and different displays for sharing and discussion [71]. In these systems, documents or data objects are shared, related and organized within a single zoomable space and each display plays a role of view to a common virtual space. Forlines et al. [72] designed a multiple display environment consisting of one tablet, one tabletop and three wall displays. They adapted two single-user, single display applications for use in their multi-user, multi-display environment, such as Google Earth and a molecular visualization tool called Jmol. In these research systems, the configuration of displays being used is already fixed by design and the main role of tabletop display is to control and coordinate the view of visualization or images from the different devices on wall displays as ancillary displays.

2.2.7 Software Frameworks for Multi-display Environments

Distributed user interfaces [73], [74] allow different elements of an application to be distributed across multiple displays. Many software frameworks for multiple displays and devices focus on supporting a distributed user interface on heterogeneous displays and devices. For example, Frosini et al. [75] presented a framework for dynamically distributing and managing UI elements of an application across multiple devices without a server, which it achieves by controlling the run-time distribution of the elements across those devices. *Panelrama* [76] is a web-based framework that allows users to distribute

the UI elements of an application into different groups called panels. The application components (panels) can be automatically reassigned to the best-fit devices based on specific device characteristics and affordances. For instance, a panel intended to display video would be assigned to a larger display, while GUI components would be shifted to a mobile display. This framework allows developers to assess the intention and importance of panel content and apportion one or more panels to the most appropriate display. Nebeling et al. [77] presented *XDStudio*, a GUI builder for enabling user interfaces on multiple displays. *XDStudio* supports two authoring modes: *simulated authoring* and *on-device authoring*. With simulated authoring, a user can design a multi-display environment on a single device by simulating target displays. On the other hand, on-device authoring enables the user to develop cross-display web interfaces directly on target displays.

The following two frameworks emphasize creating integrated visual space by both expanding visualization views and synchronizing user events across multiple devices and displays. *Munin* [78] is a framework for multi-display environments consisting of tabletops, wall displays, and mobile displays. This framework is based on a peer-to-peer architecture with three unique layers (shared state, service, and visualization layers). This architecture can minimize coupling between devices by facilitating a fault-tolerant and decentralized architecture to support a “ubiquitous analytics and visualization space.” *PolyChrome* [15] is a framework for multi-device collaborative applications that augment legacy web-based visualizations across multiple displays and manage event synchronization among them.

3 Design Considerations for Visual Analysis in Display Ecologies

Within a display ecology, the user can establish relationships between displays in a variety of ways, thereby empowering the user to accomplish different analysis goals.

Nevertheless, a display ecology presents a new set of design challenges that were not considered in the development of visual analysis tools using single displays and devices.

We argue that little research has been undertaken to address the important considerations that must be taken into account when designing display ecologies that support the analysis of data. Now that analysts are afforded increasing opportunities to conduct their analysis tasks with separate displays and devices, the design challenges associated with the analysis of large, complex data in a display ecology include these tasks:

- The demands of a given analysis task with available displays are transformed into the decision of *combining* heterogeneous displays into a holistic visual analysis environment (Figure 3.1a).

- The user needs to transfer data and information for *coordinating* information and analysis tasks across displays (Figure 3.1b).
- The user must then *connect* scattered information across separate displays in ways that enhance knowledge acquisition (Figure 3.1c).
- Throughout an analysis session, the user can dynamically change memberships of available displays in a display ecology for their changing analysis goal (Figure 3.1d).

This work primarily addresses these challenges through the design considerations for visual analysis in display ecologies (Figure 3.1). In this chapter, we distil our display ecology studies and prior research in visual analysis, information visualization, sensemaking, and human-computer interaction to extract set of design considerations for visual analysis in display ecologies. We believe these design considerations will play a critical role and provide a foundation for the design of visual analysis tools using display ecologies. This investigation will also help visualization designers, developers and researchers in understanding and evaluating new visual analysis tools using multiple displays.

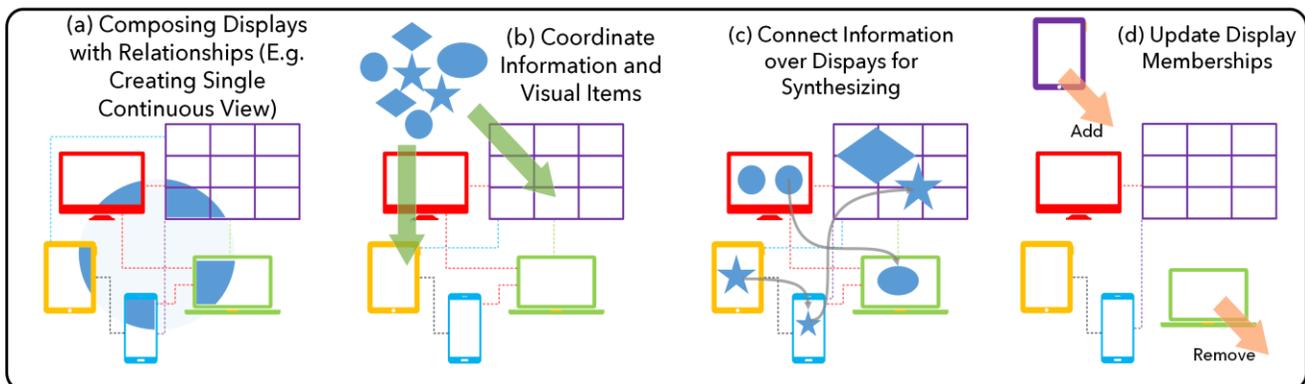


Figure 3.1. Four key design considerations for visual analytics in display ecologies.

3.1 Method

We initiated this research by exploring existing interactive systems based on multiple displays. To identify design considerations, we analyzed the literature with a focus on

multi-display systems in information visualization, visual analysis, and human-computer interaction. Specifically, we reviewed a total of 61 papers that pertained to interaction techniques, taxonomies and user studies for display ecologies and multiple display environments. The systems we reviewed via literature reports originated from a variety of communities with a number of different target applications, such as interactive media, visualization, sensemaking models, and educational platforms. Thus, these systems evidenced a fairly diverse set of design characteristics.

By prioritizing the broader data analysis activities, our design considerations were less focused on the user's cognitive/perceptual experience. Instead, our goal was to gain an insight into how users create a visual analysis space and manipulate and coordinate data through multiple displays in order to facilitate insight and knowledge formation.

To refine and group salient dimensions of the design considerations from these systems, we employed an affinity diagram method using the following steps. First, while reviewing relevant papers we wrote on a sticky note salient features and design ideas pertaining to visual analysis activities (at most a few words). As we examined and dissected more literature reports via this method, more common sets of visual analysis and interaction techniques emerged. This approach allowed us to identify various cross-display interaction and visualization techniques designed to support a relatively small set of common analysis tasks.

Ultimately, we identified a set of key insights regarding the optimal design for display ecologies. We then validated those key dimensions by revisiting and discussing whether each design consideration was truly relevant. We present in this work the various essential dimensions of display ecologies as confirmed by this validation process.

3.2 A Scenario for Visual Analysis in Display Ecologies

In general, the analysis workflow in a display ecology consists of different phases, each of which requires specific analysis and interaction techniques. This basic structure of the analysis activities in display ecologies can be illustrated with a scenario.

Let's consider, for example, an analyst wants to analyze and understand a large weather dataset with multiple displays. She first opens a multiple-view visualization consisting of three visualization views (e.g., maps, atmospheric pressure, temperature changes, etc.) on a single display; however, the resolution and size of a single display was not sufficient to show details of the multi-view visualization. Thus, she wants to enlarge each view of the multi-view visualization and decides to *combine* three available displays at her office to overcome the shortcomings of that single display's inadequate size and resolution. She then *coordinates* the three visualizations onto three separate displays in her office: a desktop screen, projector, and laptop. After gathering the multi-view visualization with different displays, she needs to understand how visual items on each display is related to others on the additional screens. She accomplishes this by *highlighting* the information across displays. In essence, she can establish connections among the visualization views on different displays by interactively selecting visual items. Finally, she can *synthesize* all of the connected information from displays in order to broaden her understanding of disconnected weather data on different displays. During her analysis, she can also *add* or *remove* different displays as the analysis progresses, instead of being limited to the current set of displays.

Although these analysis activities describe certain sequential procedures for analyzing data, we are not stipulating a strict sequential ordering of tasks and cause-and-effect constraints between phases. For example, we are not implying that connecting information is always followed by coordinating information. Also, it must be stressed that this particular workflow may not be essential for every visual analysis task within a display ecology.

3.3 Design Considerations for Display Ecologies

In this section we explore several different mechanisms by which users can link and empower multiple displays as the same visual analysis workspace. Each design consideration focuses on how users can achieve a desired analysis phase using multiple displays together. Based on the characteristics of visual analysis activities from our prior studies and visual analysis research [1], [24], [8], [21], [23], [79], we can segregate analysis tasks with multiple displays into four main design considerations (Table 3.1):

- **Display Composition:** This dimension identifies relationships among displays that are combined to create an integrated visual analysis space.
- **Information Coordination:** This dimension corresponds to how analytic tasks and data can be distributed among displays.
- **Information Connection:** This dimension corresponds to identifying a design consideration to enable users to connect and integrate scattered information across different displays.
- **Display Membership:** This dimension captures the difference between display ecologies prescribed by design and ecologies reorganized by user's adding or removing different displays during analysis.

These design considerations represent the essential analysis activities users typically need while analyzing data with a display ecology. They are not collectively exhaustive; nor are they mutually exclusive. More specialized activities may be needed, either for performance reasons or for unique analysis scenarios.

Table 3.1. Design considerations of display ecologies for visual analysis.

Design Considerations	Dimensions	Examples
Display Composition	Distributed Application Elements	Visualization elements including controls, views, user interfaces, data items, etc. across different displays are distributed onto different displays and devices.
	Data Views	More than two displays are combined to represent data views including Single Continuous Views, Coordinated Multiple Views, Navigation Metaphors, and Semantic Substrates.
Information Coordination	Synchronized	Information coordination is accomplished by synchronizing interaction and data through input redirections or a networked database.
	Surrogate	Each display offers a view into a common virtual space and users manipulate the view to coordinate information on each screen.
	Nominal	A nominal reference (e.g., a document ID, target device name, or file name or image) of both data and displays can be employed to coordinate information.
	Physical	Referring to the physical presence of displays, a user can “physically” drop information into another user’s display or any nearby display.
Information Connection	Overviews	Show the overall relationships of scattered information on separate visualization on a display; lay out and merge information items on a display
	Explicit Connection	Simple shapes such as line, triangles or circles are drawn from a source to multiple targets across displays.
	Implicit Connection	A user can organize the key person’s information on three different displays and then synthesize and form new knowledge from the structure.
Display Membership	Pre-designed	Users employ display ecologies that are fixed and prescribed by design to carry out some analysis tasks.
	Ad-hoc	Users reorganize display ecologies by dynamically adding or removing displays during analysis sessions.

3.3.1 Display Composition

Using multiple displays, users are able to construct a new alternative visual space where a greater amount of information can be visualized and interacted for enhanced analysis.

The first step in designing a display ecology is to determine how different displays can be combined and arranged to deliver the desired analysis workspace. Even though the physical arrangements of displays might be identical, the relationships among displays can be altered by specific user contexts and analysis requirements. For example, a user might want to see each view of a multiple-view visualization on a separate display.

On the other hand, another user might prefer to use multiple displays in a single logical visualization view—such as one would experience with a tiled wall display—so that the detail and size of a single visualization can expand to multiple displays.

Based on how a user forms relationships among displays for their different analysis goals and visual space needs, we classify such inter-display relationships as the following two dimensions.

3.3.1.1 Distributed Application Elements

A multi-display relationship can be considered to incorporate “*Distributed Application Elements*” when a display ecology allows users to distribute the visual analysis application’s elements—including the primary visualization view and different user interfaces—into different displays manually or automatically. Once the different elements and controls of a visualization application are distributed across multiple displays, the user’s experience with multiple displays is divided into different displays. The user’s main analysis is performed with a specific primary display, while other displays play the role of supporter for the primary display in terms of their interactions and data. Generally, the primary displays are employed to visually explore and analyze data on the primary visualization view, and the supporting displays serve as the remote controllers or data providers. For example, the visualization view of an information visualization application can be assigned to a larger display, while associated GUI widgets can be moved to a mobile display [51], [50].

3.3.1.2 Data Views

We can define a variety of visual analysis workspaces on different displays in terms of composite visualization views or structures. Thus, this relationship focuses on how integrated visualization views and structures can be constructed with multiple displays [16], rather than some relationships for application components. We can consider four different examples in this relationship for constructing data views.

Single Continuous Views: Users can combine multiple displays as a single tiled view in which the individual displays are packed together as tightly as possible to create the illusion of one single, continuous display space. An aggregated view from multiple displays increases the overall screen real estate, which enables users to see more information. Examples of this relationship include the dynamic tiling display [80], the pass-them-around prototype [81], the peer-to-peer distributed user interfaces [82], and "Junkyard Jumbotron," [56] which enables the user to combine random mobile screens into one single tiled display. Also, ScreenSquared [55] allows various web contents to be displayed over two mobile phones.

Coordinated Multiple Views: Multiple displays can also be combined for constructing coordinated multiple views (CMV), in which each display simultaneously offers alternative visualization techniques and representations for the same dataset. For example, users can assign a different visualization to each separate display. The information from those different views will ideally complement each other, thereby imparting new insights into the data analysis task. Since multiple displays are inherently discretized space, multiple displays within a single space are frequently used for CMV. One of the significant benefits of this relationship is to provide larger screen real estate for each visualization view in CMV, so users can better exploit the multiple display space to see different aspects of data.

Navigation Metaphors: A display relationship can be considered to be in "*Navigation Metaphors*" when more than two displays are combined to represent specific information visualization navigation metaphors, such as the *focus+context* and *overview+detail*

techniques. In this relationship, a large display plays the role of the main context view that shows the overview, while mobile displays provide localized views of specific data that are then positioned on that large display (i.e., the main context view). This particular relationship enables users to leverage the physical three-dimensional space around the main context display for exploring visualizations. For example, using spatially-aware mobile devices such as smartphones and tablets in conjunction with large shared displays, users are able to form a single display ecology for enhanced navigation views such as Tangible Views (Figure 3.2) [52], and iPodLoupe [53].



Figure 3.2. Tangible views and a vocabulary of physical movement of cardboard lenses [52].

Semantic Substrates: A multiple-display relationship can be considered as “*Semantic Substrates*” when the displays maintain a spatially separate, non-overlapping screen space in which different information or visual items are laid out based on related topics (Figure 3.3). Generally, this relationship does not emphasize creating the same visual structure or continuous space. The user utilizes each display and conducts analysis “in isolation and in sequence” rather than using all the displays simultaneously; in other words, the user will be switching from display to display during an analysis task. The most common example of this independent relationship is several collaborative applications running separately on different displays, such as *Design Studio* (Figure 3.3) [83], *Conductor* [24] and *VisPorter*

(Chapter 4). Nonetheless, displays in such a relationship enable users to semantically separate data across displays, as well as to distribute analysis tasks. In this scenario, a user can divide data and analysis tasks into different displays, and then perform some tasks individually on those discrete displays. However, because analysis tasks are performed within each display separately, the complexity of synthesizing distributed information is increased.



Figure 3.3. Examples for semantic substrates in display ecologies. left: affinitytable [84]. right: design studio [83].

3.3.2 Information Coordination

Multiple screen spaces facilitated by display ecologies inherently allow for dividing the analysis task into small parts based on display properties. To coordinate and divide data and analysis tasks among different displays, it is essential for a user to be able to transfer information from one display to another seamlessly. In this section, we will discuss the various trade-offs associated with the different interaction approaches for sharing and transferring information across displays. This design consideration is based on the four different approaches users employ to move information between displays.

3.3.2.1 Synchronized

A coordination technique can be identified as “Synchronized” when there is no direct data transition or interaction between two displays to coordinate different information. Instead, information coordination is accomplished by synchronizing interaction and data

through input redirections or a networked database. This approach propagates the visual analysis operations performed on one display to all the other displays automatically. Therefore, information can be shared implicitly across displays without additional interaction.

With this approach, users may perform work by having shared focus of the visualization view and simultaneous individual control of the dataset on each device [15], [85], [86]. Typical visual analysis tasks which would benefit from this approach are when each individual in the team is carrying out his analysis of data on one display but needs to have an awareness of other team members' visual analysis operations on the data. However, the introduction of unnecessary noise with the increasing number of users is one of the main drawbacks of this approach, since the display ecology shares everything instantly and indiscriminately.

3.3.2.2 Surrogate

In this approach, users can simply depend on different virtual metaphors (e.g., windows, icons, proxy, etc.) to coordinate information among displays or transfer it from one display to another indirectly. Two examples are listed below.

Windows: Sharing information can be accomplished by a large shared virtual space, and each display is represented by a localized window (which represents the view of one display) into the shared virtual space. For instance, the Zoomable Object-oriented Information Landscape (ZOIL) is a multi-display zoomable user interface framework, where each display offers a view into a common zoomable space; tangible or virtual lenses to the virtual space can be used to control or synchronize the views [70]. Thus, a user can coordinate and show visual items on different displays through indirectly moving the associate lenses in the shared virtual space on the main display without actually transferring the items from one display to another (Figure 3.3 left).

Display proxy: A display proxy allows users to visually connect to a specific display through the screen space. The display proxy provides a virtual reference for one or more displays on the other displays, which represents virtual destination display for transferring

items, as well as the availability/connectivity of different displays. For example, Wigdor et al. designed an interaction space that uses multiple wall-mounted displays and one tabletop display [42]. Each wall-mounted display features the World in Miniature (WIM) view, which is a corresponding proxy for the different displays. If a user moves one item to a different display, he or she can drag it to the corresponding WIM.

3.3.2.3 Nominal

A nominal reference (e.g., a document ID and file name or image) of both data and displays can be employed in order to share and transfer information among displays. Using nominal techniques, information sharing requires users to memorize nominal information such as URL, filename and display IDs. When users want to send or drop specific documents on a target display, they need to check the document filename and destination display's name, rather than focusing attention on the physical reference of displays or document contents. This approach is a dominant way to share files and information with others in current desktop environments. *Multibrowsing* [43] and *Conductor* are examples of this dimension. The systems force users to know the name of the desired display. The interaction techniques may hinder data analysis in display ecologies as the number of displays is increased.

3.3.2.4 Physical

The goal of the “physical” coordination technique is to enable the user to coordinate information and data items across displays physically without going through one or more such indirect and supplementary procedures (e.g., checking file names and target device names to transfer files). There are several approaches focusing on physicality and immediacy in sharing information among displays. For example, cross-device interaction techniques facilitate sharing cross-device information transfer through lightweight gestural interactions such as “flicking” and “tapping” (Section 4.2.4), or “tilting” [62] and “throwing” [87] (Figure 3.4) With these lightweight gestures, a user can “physically” throw information to another user’s display or any nearby display. In addition to the gesture-based approaches, several interaction techniques simply requires users to directly

place one device in contact with another device to transfer information between the two devices [88], [69].

A primary benefit of the seamless interaction experience is when a user needs to share a specific document, she or he can think purely of its content. Notions that may be important to provide such an experience include making the data user-focused rather than device-focused and tailoring devices to their roles in display ecologies.

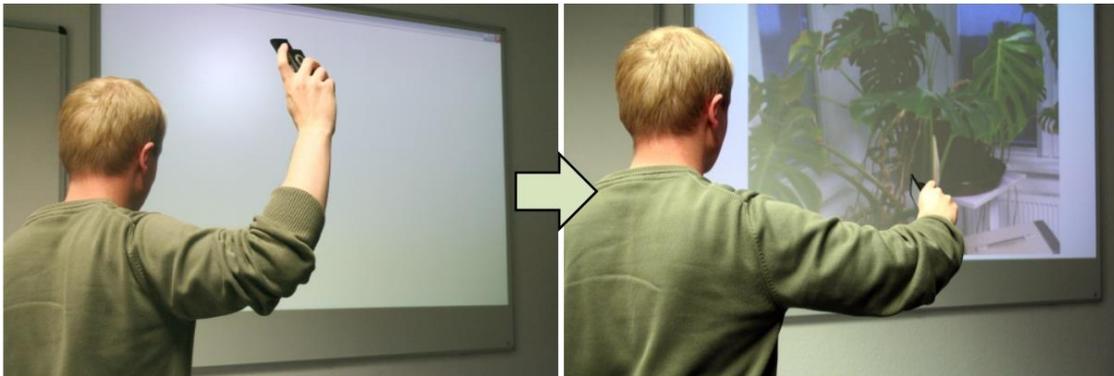


Figure 3.4. Seamless cross device interaction to move objects from one device to another – Throwing interface [87].

3.3.3 Information Connection

Information of interest and analytical activities in display ecologies are typically scattered over different displays [89]. The principal challenge associated with visual analysis in a display ecology is tied to the fact that a user must maintain awareness of and synthesize scattered information across separate displays—some of which will likely be out of the user’s immediate visual field. However, coordinated displays often require the user to switch intermittently among multiple foci of interest. Users need to connect and integrate relevant information across displays to understand data.

In this section, we explore the design considerations associated with how to connect and subsequently integrate information from different data sources (often involving different visual representations) over separate displays. Specifically, we explore visualization and interaction techniques to represent associated relationships of information items in display ecologies.

3.3.3.1 Overviews

An integrated approach can be considered an “Overview” when a display ecology offers capabilities for merging and connecting information spread over displays on a separate visualization or overlaid information items—either on a single display or via a separate view on each display. There are two main approaches (i.e., automatic and manual integration) based on this dimension.

Automatic integration: each user’s analysis activities automatically contribute to creating an overview visualization, which facilitates not only a heightened awareness of other users’ progress, but also enhances the connections between individual findings and the collective work of the group within a view. Generally, users do not actively create this visualization; rather, the system automatically creates this visualization using the information that users generate. For example, *IdeaVis* [90] presents a separate hyperbolic tree visualization view that enables users to keep track of all changes made on electronic papers and the relationships among collaborating users’ sketches on a single wall display.

Manual integration: information and visual items from different displays can be manually overlaid or merged on a single display. This approach is designed to directly integrate visualization components and views from tabletops, data walls, tablets, and laptops. This approach facilitates comparing and visually connecting data content by showing all related information side-by-side on a single screen. There are two main approaches based on this technique. The first approach is organizing and overlaying images created on different displays onto a large display simultaneously. *LivOlay* [47] and *DeskPiles* [91] system provide an easy-to-use user interface that seamlessly compares the data items (e.g., images, documents, visualizations, etc.) from different displays and devices by overlaying them on a shared display. Second, users can also merge and manipulate visualization components created from different displays and place them to compose a larger single visualization on a large display (see Section 4.2.4).

3.3.3.2 Explicit Connection

Instead of merging and connecting information and visual items onto a separate display, all information can be explicitly connected with visual representations spanning multiple displays. The visual representations represent a method for elucidating the visual relations between the data items as a single visual structure. The challenge, then, is how a visual representation might be extended in a multiple-display environment, whereby analysts can be directed to information across displays. We can consider two visual representations:

Highlighting: The simplest form to represent the connection of information across displays is to highlight one data point or a set of data points or visual items across different displays with or without some labeled text. These highlights enable users to discriminate linked data items located on different displays with different colors or shapes. The highlighting connects data items that represent the same information across different displays, thereby clarifying how those items are reflected in the different visualizations or analysis workspaces on each screen in a display ecology. In this regard, many of the current display ecology systems support information awareness and the ability to connect information on different displays via a strategy known as synchronized highlighting. The primary shortcoming of cross-display highlights is that the user must still rely solely on memory to locate and connect relevant information scattered over different displays, since this approach does not support any visual cue guiding the user's attention to relevant information across displays. This retrieval issue can become problematic when the number of highlights and displays is increased in a display ecology because users can perceive only a limited number of connections among these items on multiple displays [92], [93].

Cross-Displays Visual Representations: These cross-display visual representations can be grounded in the "partially out of the frame" approach advocated by several off-screen visualizations such as Halo [66]. A simple but effective form of cross-display visual representation is the visual link or simple shapes [94], which is one or more straight lines from a source to multiple targets across displays. Specifically, every cross-display link exits

or enters through one of the off-window directions, while a portion of that link protrudes into the off-screen area or into another user's screen space. This partial display of a visual item indicates that the remainder of that visual link resides in another display in the off-window direction. Because these cross-display visual representations are seamlessly drawn across displays (e.g., from a laptop to a tablet), they can give the illusion of one continuous workspace utilizing multiple displays while still maintaining separate workspaces on the displays. The advantages of employing one or more visual representations across displays is to enable users to connect information artifacts from different displays maintaining their analysis context on each display without switching displays.

3.3.3.3 Implicit Connection

The user can also externalize and clarify relationships of scattered information by organizing it through semantic structures and spatial relations among data, as well as displays. This implicit technique enables the analyst to leverage the multiple discretized screen space for externalization of cognitive data synthesis. With this technique, a user forms new knowledge about a complex set of data such as large document datasets. The multiple discretized screen space of a display ecology allows the user to extend the concept of "space to think" [1] by (i) by defining all display space as external memory, and (ii) by constructing semantic structures for scattered information. The user spatially organizes data within each display or across displays to transform the random layout of visual items on each display into semantically-meaningful structures (e.g., regions, people, timelines, events, or in terms of their importance to the analysis).

Some recent studies are investigating how a display ecology can support the sensemaking process [24], [70]. For example, a user can organize the key person's information on three different displays, including the telephone number (tabletop), trip route (wall), and timeline (iMac) and then synthesize and form new knowledge from these semantic structures (Section 4.4.2). These semantic structures facilitated analyst's ability to synthesize a large amount of information from the document dataset [1].

3.3.4 Display Membership

An increasingly complex digital world results in a challenge for designers of display ecologies to anticipate the varying needs of analysts. The multi-display infrastructure consists of multiple display surfaces of different form factors, including a variety of portable and semi-portable devices such as tablets and reconfigurable tabletops. Based on a wide range of individual analysis styles and needs, a designer of analysis tools can fully design display ecologies for fixed, target displays and tasks. However, the availability of different (usually large) types of datasets and the varied availability of devices may make a fully designed ecology unrealistic. In short, a growing range of complex analysis scenarios makes it difficult for designers to fulfil all the needs and possibilities with any pre-determined ecology—i.e., with a pre-designed display ecology. In contrast, it may be more effective for a display ecology to be appropriated by users in an ad-hoc manner based on different work contexts. This section addresses design issues between pre-designed ecologies and ad-hoc ecologies. The two types of ecologies can generally be determined by dynamic display memberships (Table 3.2) of display ecologies, which are characterized by the number of displays, types of displays, task division, and UI distribution.

Table 3.2. Display memberships.

	Designed Display Ecology	Ad-hoc Display Ecology
Number of displays	Pre-planned	Dynamically added / removed
Display types	Fixed display types	Unpredictable display types
Task division among displays	Prescribed by design	Divided by users
UI distribution	Different UIs fixed by display types	Consistent UIs with adjustments across displays or dynamic assignment of different UIs

3.3.4.1 Pre-Designed Display Ecologies

A display configuration can be considered a “Pre-designed display ecology” when it is designed for use with a group of target displays (i.e., a fixed set of displays). In this type of display ecology, users employ a prescribed group of displays (e.g., a wall display and multiple mobile displays) to carry out their analytical tasks. The role or tasks of different displays are fixed by design. In other words, the main goal of a pre-designed display ecology is to assign analysis tasks and data to the available devices based on functional “best fit.” A well designed display ecology will enable users to better leverage specific display characteristics and settings for analysis tasks. For example, a user can forage for information on his or her personal displays, and then multiple users can merge their information to form hypotheses on a large display. An illustration of this scenario is the Pixel-oriented Treemap for Multiple displays (Figure 3.5) [95], which is designed to divide two different visualization tasks between two types of displays for analysis of the status of 1 million online computers. For instance, each user is able to see detailed domain-specific information (e.g., machine class, function, unit, facility, etc.) on personal displays, while at the same time being able to visualize the overview of data (e.g., the overall status of computers) on the wall display. The main advantage of the designed display ecologies is that a display ecology enables users to better exploit the specific visual and analysis capabilities of different types of displays.

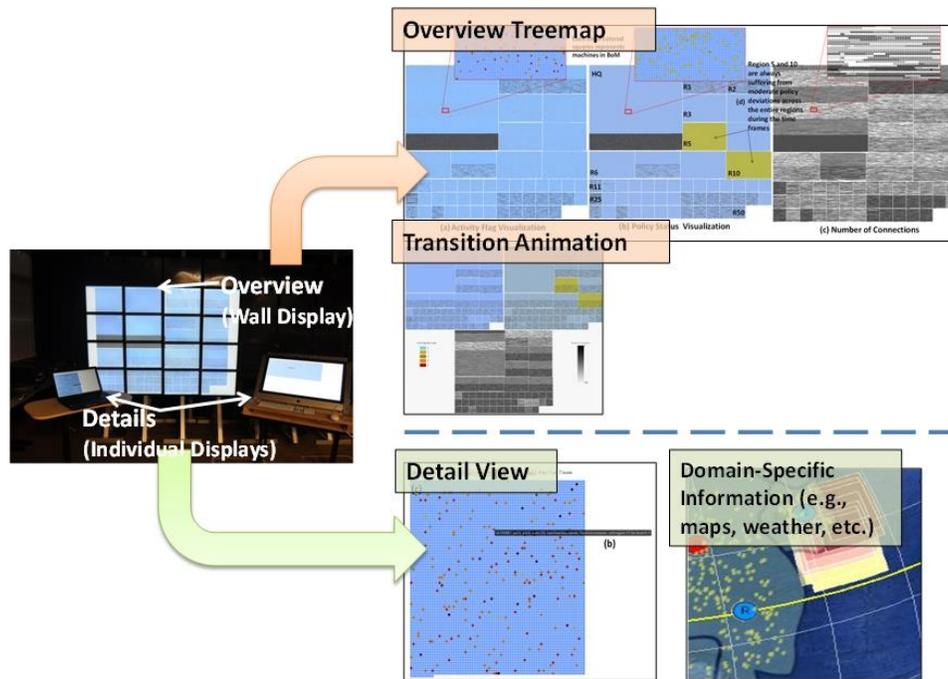


Figure 3.5. The pixel-oriented treemap [95].

3.3.4.2 Ad-hoc Display Ecologies

As noted above, the growing availability and complexity of both devices and data—coupled with the urgency of certain analysis tasks (e.g., the identification of terrorist plots)—means that analysts will be called upon increasingly to engage with diverse pieces of information and displays at opportunistic moments. Such scenarios call for the formation of an “ad-hoc display ecology,” which emphasizes the smooth reorganization and transition of available displays for different analysis activities. In this approach, a display ecology can be formed with available heterogeneous displays opportunistically. In contrast to a prescribed design for a group of target displays, this ecology focuses on creating opportunistic analysis space by dynamically assigning different tasks to and combining available displays. In this way, the user can deploy and span analytic tasks across different types of available displays in adaptable configurations and circumstances. Since analysis is not confined to a specific display, the analysis space can consist of various types of displays, including multiple large displays and mobile displays, which can even be

positioned in remote locations. Also, usable displays may join or leave the analysis space as needed.

For example, Hamilton et al. [24] presented *Conductor*, a cross-device framework that enables users to create cross-device applications by combining multiple handheld devices. With *Conductor*, a user is able to easily assign various tasks to different devices, share information, and manage different task sessions across displays through cross-device interaction methods. As noted earlier in this chapter, Rädle et al. [54] presented *HuddleLamp*, a desk lamp that facilitates spatially-aware interactions around a table by detecting and tracking the movement and position of mobile displays and the hands of users with sub-centimeter precision. This system allows for ad-hoc multi-device collaborations and interactions around a table, enabling users to mix and match different available devices. Additionally, “Phone as Pixel” [96] allows images to be drawn on the ad-hoc collection of displays.

Generally different UI elements can be shared and distributed across displays but in this ad-hoc ecology, applications on each display are designed to offer the same user experience with basic adjustments for different form factors, display size and interaction methods (touch, keyboard, mouse, etc.).

3.5 Discussion

In prior sections we explored and analyzed important design considerations for forming display ecologies for visual analysis—and in particular four crucial design aspects. Based on these design considerations and example techniques, we suggest several design advantages facilitated by the use of a display ecology. In this section, we discuss how these various approaches can further augment the design of future visual analysis tools in a display ecology.

3.5.1 Balance Foraging and Synthesis Approaches

The different design considerations suggest specific implications for analysis processes in display ecologies. While some are useful for navigating and exploring data, others focus more on facilitating analysts' cognitive analytical reasoning and sensemaking processes. We can further divide our design considerations along a spectrum of foraging-oriented and synthesis-oriented approaches in terms of visual analysis and sensemaking.

A foraging-oriented approach concentrates primarily on perceptual issues and relies heavily on the specific relationships among displays in terms of their integrated visualization views and structures—both of which are critical for exploring analytic results. Strong foraging-oriented approaches suggest specific dimensions in data views, such as “single continuous view” and “navigation metaphors.” It should also be noted that the foraging approach is concerned with gathering, verifying, and visualizing information. This approach works best when the goal is to spend a considerable portion of an analysis searching, filtering, reading, and collecting information using multiple displays. We recommend that this approach be used to increase the overall screen real-estate in order to visualize more data in geographical and multiple-view visualization applications enabled by display ecologies.

While a foraging-oriented approach focuses more on perceptual issues of data analysis within a display ecology, a synthesis-oriented approach concentrates on cognitive issues associated with synthesizing information. Specifically, a synthesis-oriented approach emphasizes externalizing the user's thought processes by organizing and distributing the collected information on a single display or multiple displays. In this approach, the physical location and presence of separate displays may play crucial roles in how to construct an analysis workspace for enhanced information synthesis. For example, by utilizing both the “semantic substrate” and “semantic structures,” users may semantically divide different displays according to types of information, the importance of information, or other task-based considerations. As both Andrews et al. [1] and Robison [21] confirmed, arranging documents into increasingly formal and meaningful structures

(i.e., spatial clustering or ordering) enable one to externalize sensemaking processes, which include data diagnostics, pattern discovery, and hypothesis formation.

However, visual analysis methodologies must sustain a broad range of analytic activities, including foraging and synthesis activities. As Vogt et al. [18] described, by supporting the specific responsibilities of these two foraging and sensemaking (i.e., synthesis) loops, one can achieve very good performance in terms of analysis for collaborative sensemaking. An important attribute of visual analysis is its flexibility in balancing both approaches to achieve a desired goal. Although there are various ways to facilitate this objective, the most powerful way to support both foraging and sensemaking is to exploit different displays.

Little is known about how these two approaches can be balanced and distributed among analysts and displays towards the goal of promoting efficiency in managing and analyzing large datasets. Hence, one important research avenue for visual analysis in display ecologies would be to investigate how to balance those approaches using different displays.

3.5.2 Exploit the Physicality of Displays

Physical space is essential for insight formation since we are embodied beings who live in the physical world [97]. In display ecologies, the physical properties of each display (e.g., its physical shape, size, specific form factors, etc.) will guide users toward which device interactions are possible, as well as how they can best be employed for a specific analysis task. We define the term, “Affordances of Interaction,” as the perceived configuration of (physical) interactions between devices that will facilitate a more natural appropriation and composition of available displays for visual analysis. Norman’s [98] concept of Affordances of products helps us to understand the optimal interaction affordance needed for the design of display ecologies. In other words, a user can exploit certain physical affordances of different displays to enhance the physical interaction between displays. An example of employing affordances for smaller devices (e.g., tablets or smartphones) could consist of directly placing the phone or tablets in contact with a tabletop to transfer

information [3] (Figure 3.6 right). Another notable example of using physical affordances in display design is the Stackable, which is a tangible widget set for faceted browsing. Each faceted token plays the role of a search parameter [2]. Specifically, each faceted token can be stacked for multiple queries such that if users want to execute a query with multiple parameters, they can simply create a stack of related stackable faceted tokens (Figure 3.6 left).

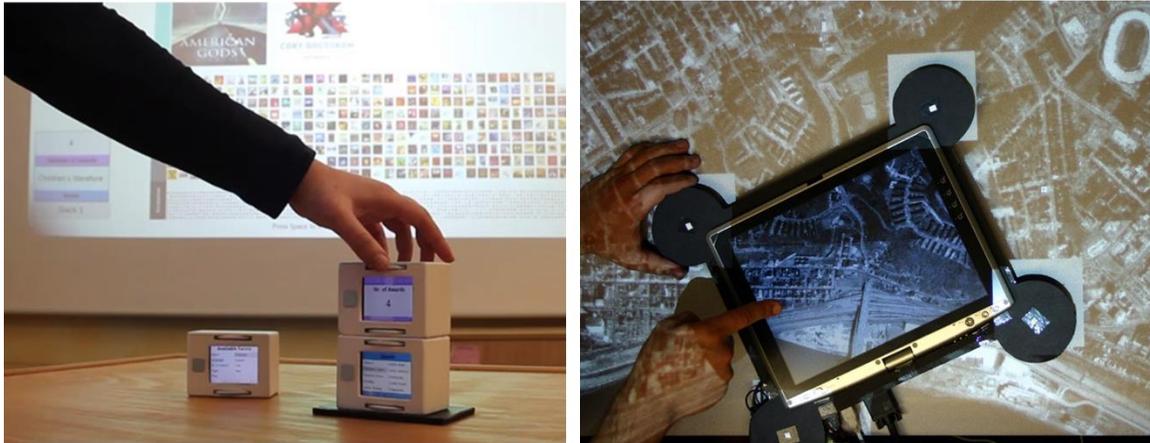


Figure 3.6. Exploit affordance of interaction for multiple displays. The stackable interface (left) [2] and moveable focus+context displays (right) [3].

In addition, the physical presence of each display provides the capability to impact insight formation. By embedding analysis components into different displays, we can create a more natural approach for analyzing big and/or complex data. For example, as described in Chapter 4, we detail a phenomenon known as the objectification of information, which facilitates the consideration of concepts on various physical displays as efficient representational proxies (Section 4.4.4).

The ways in which analysis tasks can be enhanced by the choice of and interactions with the physical properties of displays will create a more seamless environment for visual analysis. As such, we believe that further research should be conducted as to how to tap into the physicality of different displays, thereby allowing users to perceive more intuitively the possibility of cross-display interactions.

3.5.3 Provide Spatial Inter-Awareness of Displays

When users consider the spatial awareness among displays in designing analysis tools and techniques, they will be better able to leverage both the physical space and the multiple screen space afforded by a display ecology. The physical location and angle of each display will play a crucial role in how to construct an analysis workspace, as well as how to synthesize information. In general, visual analysis tools for display ecologies are designed to enable people to distribute ideas around a physical space—provided that they can be seamlessly transferred to and shared among different displays. Many cross-display interaction techniques for analysis tasks rely on the spatial reference of displays. For instance, several cross-display interaction techniques enable users to focus solely on the spatial reference of displays [59]. Additionally, spatially-aware displays can directly couple visualization environments and physical environments. Some data exploration tools allow users to customize and adapt views spatially according to the location of displays [63], [68]. Utilizing such tools, mobile displays can be relocated to achieve the desired interactions with other displays and components for enhanced visual analysis. These systems are capable of tracking the physical location of each device and detecting when they are in mutual proximity by utilizing a motion-tracking system capable of body or object tracking. These tools enable the creation of an effective spatial reference system for other displays and devices in a given physical space.

It should be noted, however, that these systems still used a simple spatial display topology, meaning that every display will be proximally located. Therefore, these potential display spatiality and topology problems suggest the need for future studies to investigate the implications and impact of selected display ecology configurations.

4 VisPorter: Facilitating Information Sharing for Collaborative Sensemaking in Displays Ecologies

Several benefits can be derived from interactive workspaces using multiple displays and devices due to their specialized characteristics. As mentioned in Chapter 2, the fact that multiple displays provide a physical space beyond one single virtual raster space enables users to (1) increasingly utilize space as a resource for visual perception and spatial ability [70], (2) with appropriate technology extend the device they are currently using to any nearby devices as needed [5], [99], (3) tap into the potential of different types of technologies for suitable tasks (e.g., enhanced data analysis) [12], [5], and (4) collaborate more flexibly through the use of multiple devices by satisfying the analytical needs of multiple users in a group [100].

These benefits are directly related to the spatial, opportunistic and collaborative nature of multi-display environments. Multiple displays enable analysts to employ and extend visual space, but require users to switch intermittently between activities and foci of interest across different displays. Thus, one of the significant inherent challenges that accompanies the use of multiple types of displays for visual analytics is the requirement for seamless cooperation and coordination of displays and devices into a unified system in which users share and subsequently integrate information and analysis tasks [89]. Although a sizable body of research describing cross-device interactions in multiple display environments is available [13], [59], [69], [101], little work has focused on directly supporting visual text analytics for collaborative sensemaking, in which multiple users can spatially and opportunistically transit and organize their analytic activities, documents, and visualization across displays.

To address these issues, we present VisPorter, a collaborative text analytics tool designed to support sensemaking in multiple display environments in an integrated and coherent manner (Figure 4.1). Through lightweight, spatially-aware gestural interactions such as “flicking” or “tapping,” the system allows multiple users to spatially organize and share both information and concept maps across displays. VisPorter provides a suite of sensemaking tools with which users can forage for information, and make sense of and synthesize it to form hypotheses collaboratively across multiple displays. We conducted an exploratory study to investigate how such a multi-display workspace, which allows users to seamlessly distribute information and visualization across multiple displays, can impact the strategy and process of collaborative sensemaking.

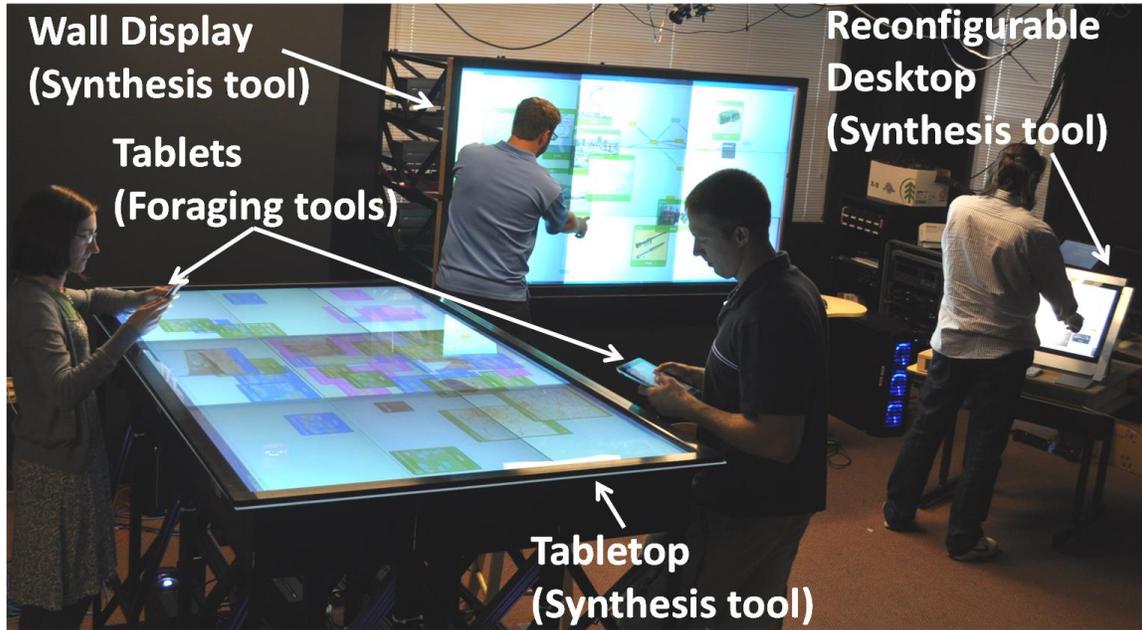


Figure 4.1. VisPorter is a collaborative text analytics tool for multiple displays.

4.1 Design Goal

Sensemaking plays a key role in the analysis and decision-making processes involved in sifting through vast amounts of information. This term can be defined as a process in which pieces of information are collected, organized, and synthesized in order to generate a productive conclusion, as well as to initiate new questions or lines of inquiry [79].

Robinson et al. [21] and Andrews et al. [1] have shown that analysts conducting sensemaking tasks with document datasets will typically utilize a large physical space (i.e., table and large display space respectively) to externalize their thought processes by spatially organizing the documents—in essence by defining the space as external memory. To enable such external memory and semantic structure with display ecologies, the most important design requirement is how users can share and coordinate information among different displays. We need to consider some natural interaction methods that enable users to spatially arrange analytic tasks and data over displays, as well as to help them immediately decide what information can be shown on different displays.

Guided by the design consideration (Chapter 3)—and coupled with findings from several prior related research projects in visual analytics, sensemaking, large displays, and multiple display environments—we generated four design principles (D1-D4):

D1. Exploit physical space through physical navigation and persistence:

Physical space is essential in sensemaking since we are embodied beings who live in the physical, tangible world [97]. For example, Ball et al. [17] demonstrated how physical navigation produced a performance improvement in visualization tasks over virtual navigation. They proposed several design suggestions to facilitate physical navigation in the design of visualization systems, thereby reducing dependency on virtual navigation (e.g., scrolling, panning, zooming, etc.).

D2. Share visual information objects in a direct and physical manner:

Generally, access and management of dispersed information across multiple devices is a major problem in multiple display environments. For an integrated multi-device system, users must be able to share and analyze information objects and visualizations in a direct and intuitive manner. Moreover, the user should be able to focus attention on the direct physical reference of the material being handled (e.g., a particular document, entity, and image), rather than relying on the nominal reference, such as a document ID, filename, or URL. Nacenta et al. also confirmed that the ability to maintain focus on the material being handled during spatially-aware interactions is preferred for transferring data between devices [59]. Chu et al. identified five design themes that relate to how multiple devices may help an individual's thinking processes by physically objectifying information [102].

D3. Spread and organize tasks, data, and visualization across displays:

Devices should independently allow for the maintenance of data, workspaces and analysis activities based on display form factors, while ensuring that the end results of personal analyses and data sources are incorporated into the final unified results. For instance, a multi-device system should facilitate both individual analysis and synthesis tasks, as well

as seamless transitioning between tasks. Vogt et al. provided several design suggestions for co-located collaborative sensemaking using a large shared display, and found that collaborators frequently preferred different analytic approaches, sometimes requiring different devices [18]. Geyer et al. also suggested that different activities such as individual or collaborative tasks should be supported by suitable devices and modalities [100].

D4. Support dynamic device membership and spatial inter-awareness:

Users should be able to easily reorganize analytic workspaces across displays based on changing needs, and to deploy and span analytic tasks across the different types of available displays. Therefore, the necessity to interrelate devices and user activities implies that an interoperable infrastructure supporting dynamic display membership in multi-display environments is a must. Such a system can be supported through a plug-and-play model that enables the user to pick up, mix and match displays, tasks, and interaction techniques. With such an infrastructure, all displays enable continuous support and capture the insight formation process as it occurs in any display or over time in a larger information space.

The above design principles, derived from the literature and insights from our own past sensemaking research projects, formed the foundation of our design choices during the development of VisPorter. We reference the principles throughout the remainder of the chapter to describe the system itself and how these design principles supported, hindered or modified users' behaviors with the system during the study.

4.2 The VisPorter System Overview

VisPorter was designed with the goal of achieving collaborative insight into a large number of text documents by sharing, transferring, and spatially organizing digital objects in multiple format types and multiple visualizations across displays. It also supports synchronous, collaborative creation of concept maps from a set of important keywords

across different displays. In the following sections, we illustrate how we designed the tools and interfaces of VisPorter through a use case scenario and then we describe the tools and important capabilities of VisPorter in greater detail.

4.2.1 Usage Scenario

We consider two analysts (“Ava” and “Ben”) who are collaborating on the investigation of a large dataset containing 1700 documents, including intelligence reports, news articles, web pages and pictures, in order to uncover the hidden story (such as the VAST Challenge 2007 scenario [103]). The two analysts use the VisPorter System on two tablets individually, and share a touch-enabled tabletop and one large wall display.

Both Ava and Ben start their analyses simultaneously using the Foraging tool on their personal tablets (Figure 4.2 & Figure 4.3) independently. They quickly read many documents on the Document viewer (Figure 4.2) in order to familiarize themselves with the data and find potential key persons or other keywords that appear repeatedly. Based on these key entities, each analyst performs searches (Figure 4.2a), reads associated documents more carefully (Figure 4.2d). Ava first focuses on the automatically highlighted entities on the Document viewer (Figure 4.2d) since she can see the entity type by color; however, she finds that there are keywords and unknown names that are not identified and highlighted by the system, so she adds them as new entities. If new relationships between specific entities are identified while reading a document, Ava and Ben establish connections between two related entities. For example, Ava adds a relationship between the “Sally” and “tropical fish” concepts and labels it “is a marketer of.” Ava verifies and removes some incorrect relationships between entities for the current document (Figure 4.2e). The analysts also begin bookmarking the interesting documents or throwing them to the large displays or to the other user’s tablet.

However, as their individual analyses progress, both analysts encounter difficulties in sharing their findings or important insights due to the physical separation of their individual lines of investigation on each tablet—which means they lack direct awareness of what the other analyst is working on. Thus, they decide to directly share and collect

documents, pictures, and concept maps on the wall and tabletop displays (Figure 4.4 & Figure 4.5). Both analysts flick the documents in the direction of different displays on the document viewer when they find interesting information or want to reference them later and tap important entities to share the concept map with another analyst (D2). Viewing shared documents on the common space facilitates the direct sharing of interesting pieces of information and discussion about their immediate findings. For instance, while the analysts discuss an epidemic outbreak, Ben wants to know when the outbreak was first noticed. Ava immediately flicks the document related to the time line of the outbreak toward the wall display for Ben to observe (D2).

As the number of documents on the shared display increases, Ben wants to better understand the relationships of the collected documents on each large display. Therefore, they start organizing documents spatially on the wall display and tabletop using various central factors, such as locations and timelines (D1, D3).

The analysts build the concept maps collaboratively as they continue identifying and establishing relationships between entities. As the investigation progresses, Ben wants to see a larger concept map that includes more entities, but it is difficult for him to see all related entities on the small screen of the tablet. So he visualizes the larger concept map on the wall display by selecting and tap-holding multiple entities on the ConceptMap viewer (Figure 4.3) to transfer them to the wall display (D2).

They move between two large displays to analyze shared information and to discuss questions about documents organized on different displays (D1, D3). They often refer to their tablets for individual analyses. The spatial organization of documents across displays (D1, D3) facilitates convergence to a common understanding of the results. In short, the two analysts successfully reached a common hypothesis using VisPorter.

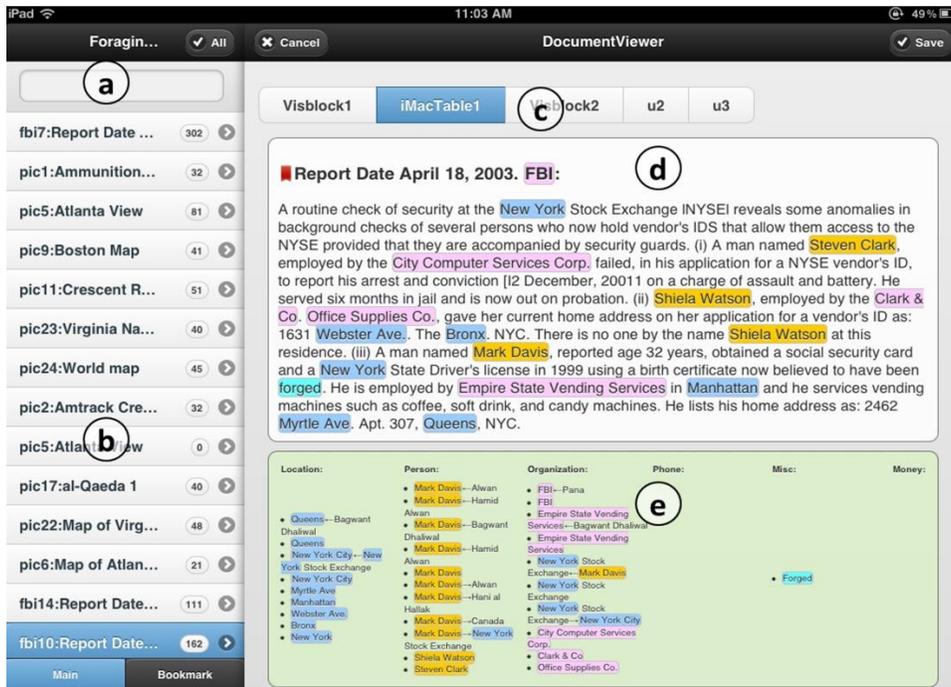


Figure 4.2. Foraging tool - Document viewer.

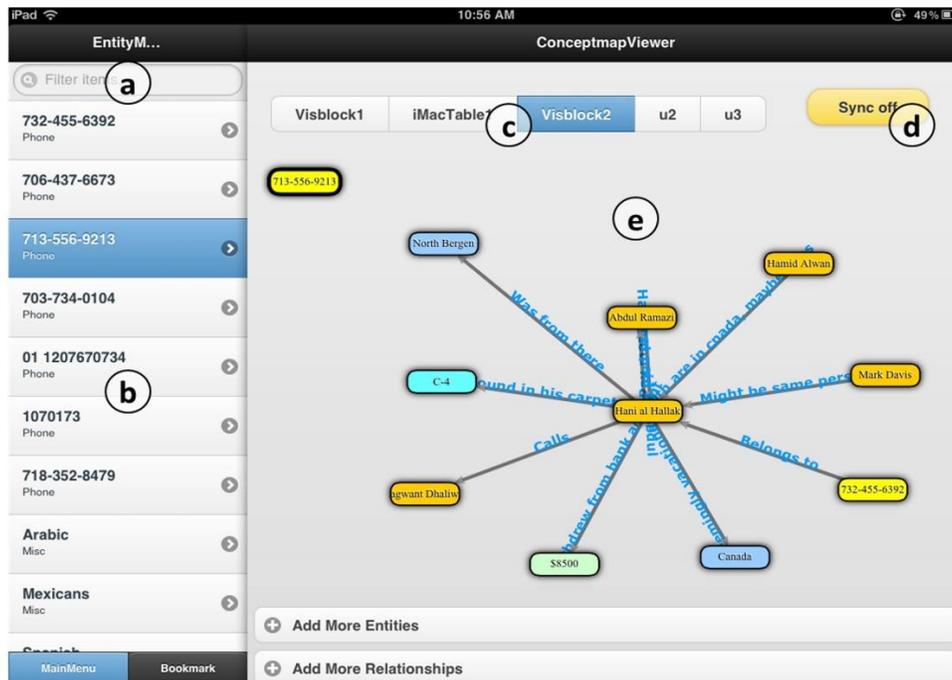
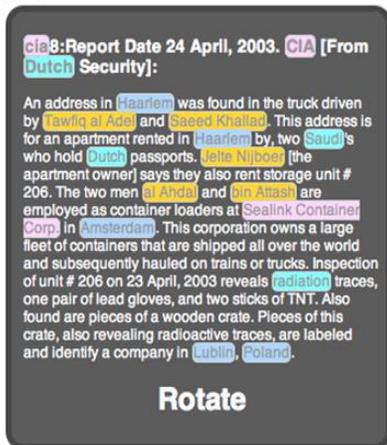


Figure 4.3. Foraging tool - ConceptMap viewer.

(A)



(B)



Figure 4.4. Two types of document boxes for the synthesis tool (a) text document and (b) image.

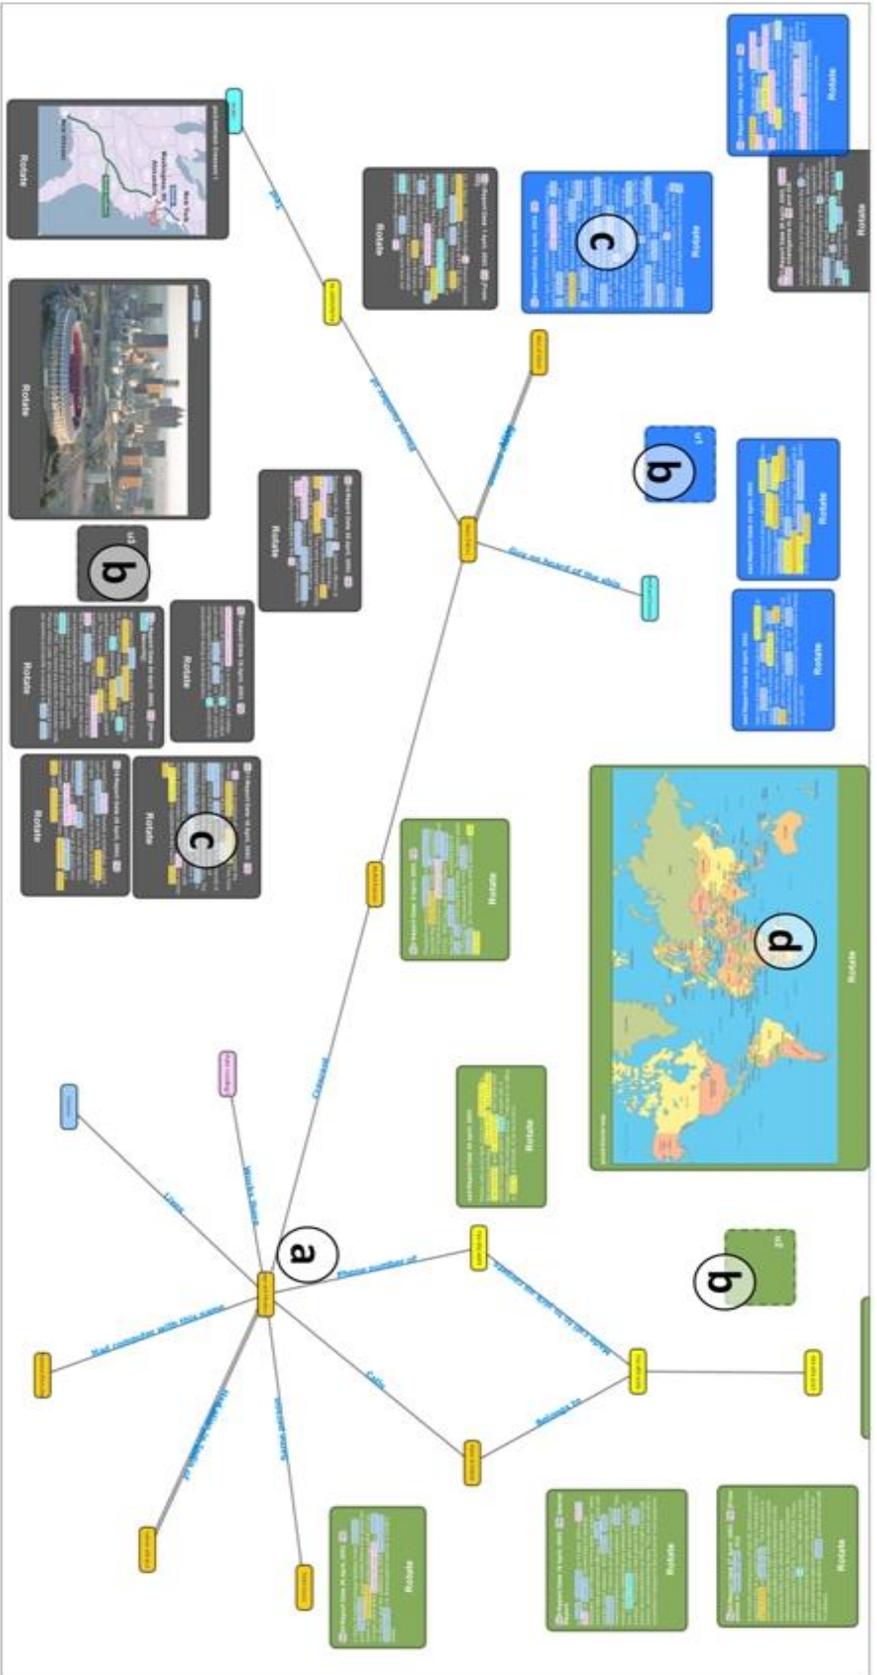


Figure 4.5. Synthesis tool on the shared display.

4.2.2 Sensemaking Tools

The VisPorter system consists of two main sensemaking tools: the *Foraging tool* (consisting of the *Document viewer* and *ConceptMap viewer*) and the *Synthesis tool*. Each of these is primarily designed to support different stages of the sensemaking process [79]. These two tools directly match the two sensemaking loops in the Pirolli and Card model: the Foraging and Sensemaking Loops, respectively. As Vogt et al. confirmed, supporting the division of responsibilities for these two loops showed very good performance in analytical tasks requiring collaborative sensemaking [18]. One way to achieve this is to utilize two specialized tools for foraging and sensemaking, which are supported by a suitable display affordance (D3). The user interfaces for the Foraging tool are designed for personal analysis and devices easily carried by users, such as tablets and smartphones (Figure 4.2 & Figure 4.3). The Synthesis tool allows users to take advantage of large screens by organizing documents and concept maps spatially on the screen, as well as by enabling the integration of various data from multiple users and devices (Figure 4.5).

Foraging Tool:

The Foraging tool facilitates sorting data to distinguish what is relevant from the rest of the information. The individual spaces provided by the foraging tool were inspired by the foraging loops of the Pirolli and Card's sensemaking model [79]. Even though users are collaborating on the analysis, they need to spend a considerable portion of their work searching, filtering, reading, and collecting relevant information individually [18]. This tool is designed to facilitate these individual tasks on personal devices. The tool includes two main viewers – the *Document viewer* and the *ConceptMap viewer*.

- *The Document viewer* focuses primarily on individual content exploration and identification of important entities and their relationships (Figure 4.2). Discretized foraging space is useful for user's sensemaking tasks. Users can read, search, retrieve and bookmark raw data such as text, images, etc. via a mobile application interface. The viewer allows multiple keyword searches, as well as the creation of entities, relationships, and annotations for each document. A search result is ranked and ordered by tf-idf [104] values for the keywords. The viewer includes a document

(Figure 4.2d) and an entity-relationship list (Figure 4.2e). Users can add entity or relationship interfaces and annotations through the similar interfaces used in VizCept [8]. Each document is automatically parsed for entities using the LingPipe library [105] and the extracted entities are highlighted in different colors based on entity type (e.g., people, locations, etc.). At the top of the interface (Figure 4.2c), toggle buttons show a list of target devices that can communicate with the device in use; these buttons are dynamically updated based on available displays.

- *The ConceptMap viewer* allows users to visualize entities and relationships in a force-directed layout concept map [106] (Figure 4.3). Users can add, select, remove, and search within the created concepts on the entity list panel (Figure 4.3b). In the right panel, selected concepts from the entity list panel are visualized in the ConceptMap viewer. A user can drag and drop entities or concepts onto the ConceptMap viewer using touch inputs. Like the Foraging tool, the ConceptMap viewer allows users to create entities and relationships via the collapsible user interface or by simply tapping specific entities (Figure 4.6). The viewer has a Sync button (Figure 4.3d), which when switched on directly shows the personal controls and views of individual concept maps on the Synthesis tool of the target large display.

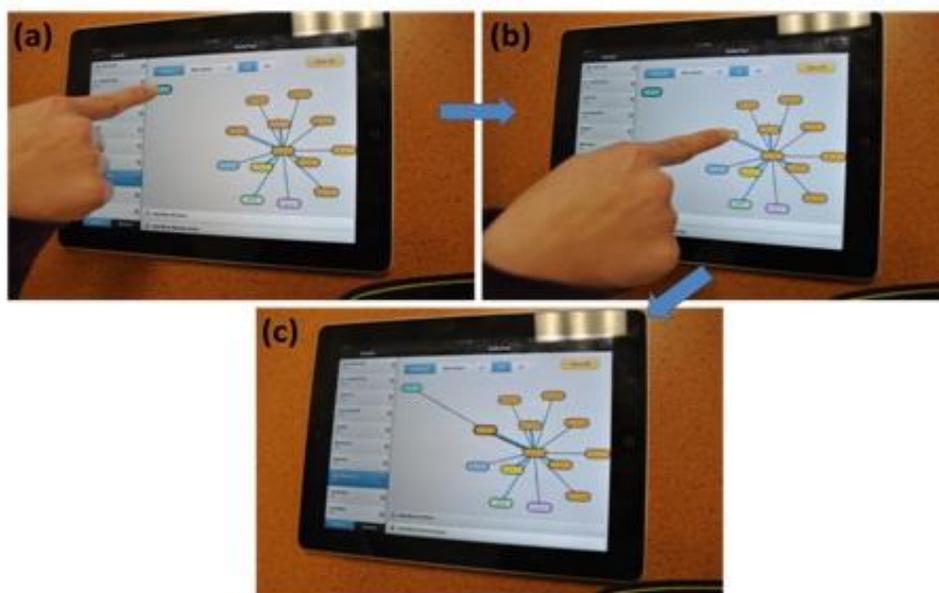


Figure 4.6. Easy to connect between two entity nodes by tapping gestures.

Synthesis Tool:

The Synthesis tool involves utilizing the information pulled aside during the foraging process to schematize and form a hypothesis during the analysis. This tool emphasizes collaborative synthesizing of the collected information on the shared space, while the Foraging tool is concerned more with gathering, and verifying information. The Synthesis tool enables the user to integrate findings that have been collected on different devices by dragging and dropping information (e.g., documents, images, concept maps, entities) (Figure 4.4 & Figure 4.5). Figure 4.5 shows documents (Figure 4.5c) and a concept map (Figure 4.5a) created by users. The Synthesis tool facilitates spatial organization of the *information objects*, which include text documents (Figure 4.4a & Figure 4.5c) and images (Figure 4.4b & Figure 4.5d) from different users and different devices (D1, D3). As with the Document viewer in the Foraging tool, entities are highlighted in the Synthesis tool.

4.2.3 Display Proxy Interface

In the space created by VisPorter, users and portable devices need to move around another display and users often need to transfer documents from one display to a specific location on a nearby display. So, moving an information object between devices relies on the physical presence of devices and their locations. To show other displays' physical locations, VisPorter provides an interface "Display proxy" which allows users to spatially and visually connect to a specific device through the screen space (Figure 4.5b). When a new device engages one of the VisPorter tools, all other devices display a visual reference to the associated display proxy on the Synthesis tool. The display proxy provides a spatial reference for the specific display on the other displays. It represents spatial target positions for transferring objects as well as the availability/connectivity of different displays.

The proxy is designed to support motion-tracking systems which enable devices to detect when they are in mutual proximity. If the proxy is connected to a motion-tracking system capable of body or object tracking (e.g., VICON, Optitrack, etc.), it is an effective spatial

reference for other displays and devices in a given physical space. If motion tracking is not supported, these proxies can be dragged and dropped on the screen space for users to manually determine a drop position.

4.2.4 Gesture-based Interaction

In VisPorter, users can “physically” throw a piece of information to someone who is nearby or to a large screen with the flick or tap of a finger through the use of two different types of VisPorter tools (D2). All information objects including text documents, images, and concept maps are transferred around the location of the display proxy on other large displays. VisPorter employs gesture-based techniques for moving an information object between the Foraging tool and Synthesis tool. When users transfer an information object from the Foraging tool to the Synthesis tool, the position where it is dropped can be determined by one of the four swiping directions (i.e., up, down, left and right) (Figure 4.7). For example, if a user swipes toward the right side of her tablet, then the flicked document is dropped on the right side of the associated proxy on a target large display.

The tap-hold gesture is also used to transfer an entity or concept map to the Synthesis tool, and users can merge individual concept maps with the larger concept map on the Synthesis tool through the tap-hold gestures (Figure 4.8). For instance, multiple users can create their individual concept maps independently on personal displays, and then combine them into a large concept map on a shared display (e.g., wall or tabletop displays). Generally the size of the entities on the screen is fairly small, so tapping is a more useful gesture than swiping to transfer the concept map (entities).

On the other hand, moving documents or entities between two large displays running the Synthesis tool is carried out through display proxies and simple gestural interactions. If a user wants to send a copy of a specific document from the synthesis tool on a tabletop to a wall display, she can simply tap-hold both the document and a display proxy of the target display at the same time.

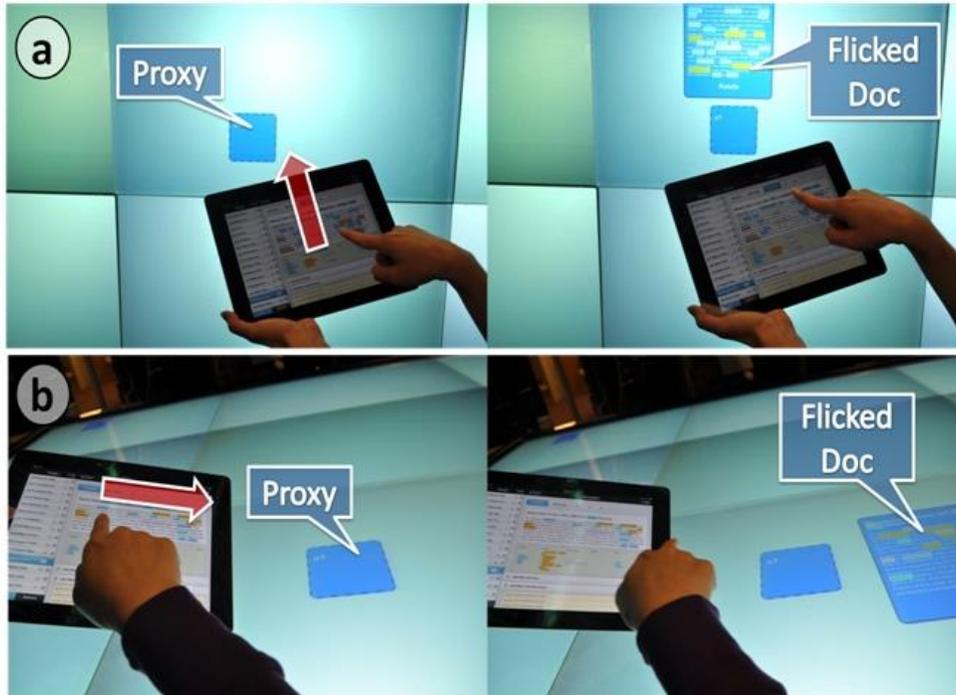


Figure 4.7. Swipe and drop the document onto the shared displays: (a) Wall displays and (b) Tabletop display.

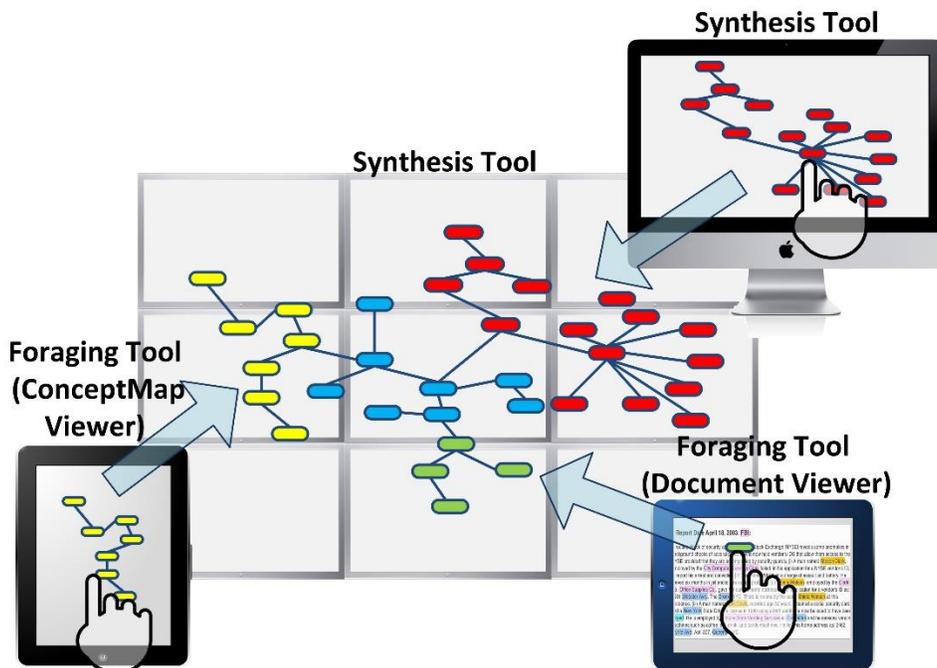


Figure 4.8. Transfer and merge individual concept maps and entities in a wall display through tap-holding gestures.

4.2.5 Implementation

To support interoperability and spatial inter-awareness (D4) among different types of devices, we employed a web architecture for VisPorter, which consists of multiple web clients and a server. This architecture is based on bidirectional communications among multiple devices and applications via Websocket [107], which enables a persistent socket connection through a server. In our infrastructure, the data (e.g., user gesture events, documents, entities, concept map data, etc.) between the client and server are exchanged in compressed Java Script Object Notation (JSON) format [108].

To ensure support for interoperability, an important issue is how the information produced by different displays is distributed and synchronized. The clients provide user interfaces and visualization views in which information objects and concept maps are displayed. All clients (devices), such as the Foraging tools and Synthesis tool, are independent web applications that share application state information, input events, data queries, etc. with other clients through the server. All communication between devices (clients) is mediated by the server. For example, when a gesture event (i.e., flicking a document) occurs on a client on a tablet, an associated message comprised of gesture types, information queries, target device id, user id, and document id in JSON is sent to the server. The server then processes the JSON message by retrieving a flicked document from the database and returning requested documents to another client on a target device. The server also keeps track of device configurations and the status of applications in order to manage distributed software and hardware resources in VisPorter. To manage the location information of each hand-held device from a motion tracker system, the VisPorter system maintains an independent input server, which transmits each device's location information to the server.

VisPorter clients (i.e. the Foraging and Synthesis tools) are implemented with JavaScript, HTML5, CSS and JQuery (for the foraging and entity tool) and the servers are implemented with Node.js [109]. To use the touch interfaces on the wall and tabletop displays, we used TUIO [110]. The concept map is developed with HTML5 Canvas.

Since the information objects are based on a form of DOM elements, users can wrap various common data types (such as text, images and videos) and various web services in the DOM elements.

4.3 Evaluation

We conducted an exploratory study of our VisPorter system using various types of touch-enabled displays. We had two main goals. The first goal was to better understand how the multi-display environment created by VisPorter impacts the users' processes of co-located collaborative text analytics. Specifically, we wanted to extend previous findings [1] about how users conducted analysis tasks on single large displays to examine how users externalize their synthesis activities into the physical space provided by a multi-display environment. Thus, this study focused on investigating how users employ multiple display spaces to collaboratively create semantic structures over multiple displays, as well as to utilize the discretized screen space as external memory for information foraging. The second goal was to evaluate how well the design (D1-D4) appropriately supports the sensemaking tasks in collaboratively solving complex problems with our tools and multiple displays.

4.3.1 Participants

We recruited 24 participants, 4 females and 20 males, from a pool of computer science graduate students whose ages ranged from 20 to 39. Our sample reflected the existing male-to-female ratio in the computer science department from which the participants were recruited. A pre-session survey confirmed that none of the participants reported familiarity with the use of large displays or tabletop displays. All participants were required to have prior experience with visual analytics or information visualization by having taken a course on either topic. While they were not actual analysts, they had basic knowledge about how to approach analytic problems from their required graduate-level classes. Prior user studies in collaborative visual analytics have also made use of participants without formal training as data analysts [41], [18]. The participants were

grouped into eight teams with three members each (G1 to G8). Four teams included members who knew each other beforehand, but the other four teams did not (Table 4.1).

4.3.2 Task

While sensemaking occurs in many domains, in this work we focus on document analysis. In this study, users performed an intelligence analysis task, in which they analyzed a collection of intelligence documents to identify potential illegal human activity and motivation. Each team conducted the analysis in a co-located synchronous fashion using VisPorter in a multi-display environment. The task, which did not require any specialized knowledge, was to identify a latent plot hidden within a fictional intelligence dataset [111]. The dataset consisted of 41 documents and 25 pictures, and included three subplots that comprised the main terrorist plot. The dataset was relatively short and of an appropriate size to complete within the one-hour time limit, as in prior work [21], [18]. The task also included “noise,” with the potential to lead users to unrelated hypotheses.

Participants were asked to use VisPorter to forage information from the dataset that most efficiently led to productive hypotheses, and then to synthesize information from multiple intelligence reports. Their goal was to provide a hypothesis solution with supporting evidence including details such as who, what, where, when, and how these pieces of evidence were connected. Before starting the analysis, all teams were given an answer sheet to complete during the task. This answer sheet asked the teams to provide short answers to four questions based on [112], including the entire situation and plot, key persons, the timeframe of the plot, and the important locations of the plot. The short answers were graded by an author, as shown in Table 4.1. The grader awarded each correct answer 1 point. The maximum possible score was 10 points.

4.3.3 Apparatus

A suite of devices comprised of iPads (one for each participant), a touch-enabled iMac (with a tilting screen to allow for tabletop or vertical use), a shared wall display, and a tabletop display were made available to the participant teams during the study. These displays provided very different affordances. The eight teams had access to all devices at

all times during the analysis and the participants were free to choose devices based on their needs. Both the tabletop and wall display were made of nine tiled-back-projection displays arranged as a large 4ft by 6ft (3840x2160, 82.5 inch diagonal) horizontal or vertical surface screen with a PQ Labs' 32-points Multi-touch overlay.

4.3.4 Procedures

The study was carried out with each of the eight teams conducting a 1- to 1.5-hour-long analysis session in a laboratory environment. All three team members met in the lab at a scheduled time. A demographics questionnaire was administered to each participant and then they all underwent a 20-minute training session as a group on how to use the system. The experimenter first provided a brief demonstration and explained the two main tools of VisPorter; he also introduced the set of available displays and devices. During this training session, users could freely test each feature of the system on the different displays. However, no analytic approaches or strategies were discussed during the training session to avoid influencing the participants on their analytic tasks.

After the tutorial session, all participants started a one-hour analysis task sitting or standing in front of the large displays. The dataset was preloaded before the study and the questions were then shown. The Foraging tool was activated on the iPads and the Synthesis tool was started on the wall, tabletop and iMac displays. During the analysis, participants were allowed to ask the experimenter how to use VisPorter.

After 1 to 1.5 hours of the analytic session, a debriefing followed, during which the participants were allowed to access their analysis results on the displays. Each team was then asked to complete an answer sheet and a post-questionnaire concerning their findings and their user experiences in completing the analysis task with the system. A semi-structured group interview was conducted at the end of the session involving all team members.

4.3.5 Data Collection and Analysis

All sessions were video-recorded and a researcher who remained in the experiment room took observation notes. Screen activity was recorded for all work done using the Synthesis

tool on the wall, tabletop and iMac displays; screenshots were taken at 30-second intervals. All concepts, relationships and notes created by the teams were logged in a database and retained. Additionally, all interview results and conversations during the collaborative analysis sessions were audio recorded and transcribed by the authors. Our analysis was mostly qualitative in nature. We analyzed the data using a grounded theory approach. An open-coding session was first performed on our collated observation notes, interview transcripts, and post-questionnaire results to uncover high-level themes—for example, the participants’ use of the various devices and their strategies for sensemaking and collaboration. The authors discussed these issues, and collated them on the whiteboard. Based on this information, we defined a set of high-level themes regarding the sensemaking process.

We then implemented a second round of more detailed coding using the high-level themes as categories. After important analytic strategies were derived, we consolidated our findings by conducting a validation procedure of those strategies by examining other types of relevant data, including screenshots, video and audio recordings of the sessions. In this section, we present the common strategies with supporting details from different sources wherever appropriate.

Table 4.1. Study result.

	G1	G2	G3	G4	G5	G6	G7	G8
Primary style	FS	SD	FS	FS	SD	FS	SI	SD
Secondary style	No	FS	No	SD	No	SD	SD	FS
Cross device organization styles	None	Single entity (location)	None	None	Single entity (location)	Multiple entities (geography/ organization/ people)	Single entity (location)	Multiple entities (telephone + money/ location + event/ people)
Score (out of 10)	2	6	3	7	4	6	4.5	8
Number of flicked documents (iPads to large displays)	11	21	22	43	16	67	50	35
Number of flicked entities (iPads to large displays)	12	20	5	37	32	1	0	102
Number of large displays used	1 Tabletop	2 Tabletop and wall	1 Tabletop	2 Tabletop and iMac	2 Tabletop and wall	3	3	3
Objectification behaviors	No	Yes	No	No	No	Yes	No	Yes
Knew each other before the study	Yes	Yes	No	Yes	No	Yes	No	No

4.4 Findings

The key results from participants' use of VisPorter, which we elucidated from our study, are summarized in Table 4.1. The table shows how many groups fell into each collaboration style, how much each team exchanged or transferred information across different devices, scores based on the identified plots, etc. It is important to stress that we did not focus on the statistical analysis of results. Instead, we are more interested in how the process of sensemaking was influenced by using VisPorter. As Huang et al. [5] emphasized in their display ecology study, our evaluation focused on how the display ecology, created by VisPorter, was able to support collaborative text analytic tasks, rather than measuring the use of VisPorter's features and displays. Each finding will relate to qualitative results and discussions described in the subsections. In our study, we observed four common strategies that the participants used during collaborative sensemaking with VisPorter.

4.4.1 Collaboration Styles with Multiple Displays

We first focused on understanding how teams worked together and coordinated their analysis tasks across the different displays. From our observations, although the participant teams had varied work division approaches, their approaches can be generalized into three types (Figure 4.9).

- **Strictly individualized (SI).** For this type, each participant had strong ownership of a specific large display in the environment (Figure 4.9 left). The tabletop, wall and iMac displays were divided among the three team members, and were used as individual workspaces in addition to the individual iPads. In this approach, the teams assigned portions of the initial information to the team members and each team member focused on individual analysis on a different large display. Members occasionally looked at the other members' displays, but there was almost no discussion or other significant collaboration among the participants during the analytic session. Therefore, until the debriefing session, these participants did not

combine and synthesize individual findings from each display. Instead, all users commented they wanted to concentrate on their individual analysis.

- **Semi-divided (SD).** Like the “strictly individualized” case, each participant had ownership of a specific large display and concentrated on working on that display (Figure 4.9 right). The team members divided the given data between the shared displays. Each member mainly worked with his or her large display. However, during the session, they looked at each display together, and shared the knowledge/insights gained from the data as needed. They often shared the findings with each other and asked their team members to come closer to the display for assistance. Once a member found possibly useful and interesting information for another participant, he/she approached that user’s display and flicked the document. However, each member still focused on an individual analysis with one display.
- **Fully shared (FS).** In this case, participants did not have specific ownership of any large display (Figure 4.9 middle). If the team used multiple large displays (G4, G6), they first discussed the categories of data and assigned each to a suitable large display based on the contents and entities. In contrast to “semi-divided,” all users spent a fair amount of time analyzing data around the tabletop display instead of each member working on a specific topic with separate displays. They shared all information with each other and collaborated to reach the goal. When they needed to organize or forage information on different displays, they immediately moved to that particular display or transferred related information from their iPads or tabletop to the corresponding displays.

Table 4.1 shows which collaboration styles each team used most often, and the second row shows other styles that they sometimes used. Four of the eight teams (50%) primarily used “fully shared,” which was utilized the most among the teams; conversely, the “strictly individualized’ approach was used least. We observed that G2, G4, G6, G7, and G8 changed to secondary styles as necessary.



Figure 4.9. Three collaboration styles for multiple displays. Blue arrows indicate users.

4.4.2 Cross-Display Semantic Structures

An important research question in our study concerned how users spatially organized and distributed their data and findings on multiple displays (D3). The discretized screen space supported by VisPorter allows users to arrange documents and entities onto different displays. We examined how analysts leveraged such discretized screen space of multiple displays to augment the information with synthesized meaning. The displays enabled the participants to spatially organize hypotheses and evidence, providing external representations and spatial insight (D1). These activities can be classified according to the evidence marshaling and schematizing stages in Pirolli and Card’s sensemaking model [79].

We observed a variety of spatial organization methods performed by the participants during their analysis using VisPorter. Spatial organization strategies of documents on each single large display echo results of previous studies on large displays [1]. For instance, the participants created spatial structures such as document clustering and ordering on the display. We also observed “incremental formalism” [113]. Some of the teams that used SD and FS styles incrementally morphed their organization of data across displays into more accurate arrangements as their analysis progressed. In this section, we focus on salient organizational strategies used with multiple large displays. The multiple display types allowed the participants to organize the data based on the device capabilities and visualization need. We observed two categories of cross-device spatial organization.

- **Single entity types:** Three of the eight teams preferred to collect information based on the geographical area of interest. We attributed this to the fact that the dataset included a large amount of location information. Thus, the teams organized data according to a single entity type—location. For instance, when G7 decided to organize the given data into three primary locational areas of interest (Virginia, New York, and Boston), each area was then mapped to a particular display—Virginia data to the wall display, New York data to the tabletop, and Boston data to the iMac. Since there were many documents related to Virginia that included locational data, they decided to use the large wall display for that data.
- **Multiple types of entities and visual representations:** Two teams focused more on arranging data in different displays based on multiple entity types. G6 organized information by different entity types such as places, organizations, people, and events in each large display. G8 also distributed data to three different displays based on (1) telephone numbers and money, (2) locations and events, and (3) people. This strategy allowed the team to use different visual representations on different displays based on the type of information being visualized. For instance, G8 formed hypotheses on three displays (Figure 4.10), based on an event timeline (iMac), people’s locations and trip routes (Wall), and telephone and bank accounts (Tabletop). On the tabletop, a concept map was presented to determine how people were related to each other, based on telephone numbers and bank accounts. Tracking the telephone numbers and money required seeing the relationships among people. On the wall display, the team opened a large map and overlapped related documents for the locations of different terrorist plots. On the iMac, participants spatially organized a time sequence (horizontally) with the anticipated travelling movements of the key people. By integrating with the location of explosives, they deduced the possible target locations.

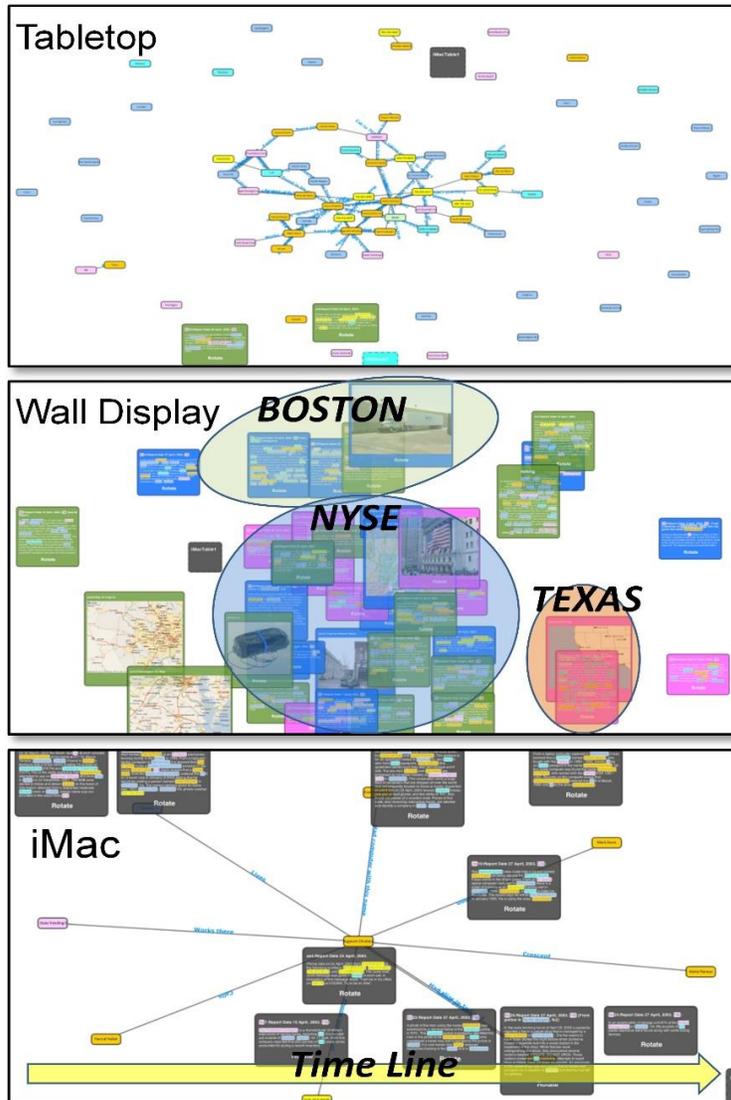


Figure 4.10. Organizing information based on multiple entity types on different displays. On the figure of the wall display, we added labels pertaining to participant explained regions of clustered documents described to us during the debriefing.

4.4.3 On-demand Extension of Display Space

We analyzed when participants “threw” information to another device and the rationale for why they transferred their activities to the chosen device. During the post-interview, all participants were asked what information and why they transferred from their personal tablet to the other displays.

Offloading Information. We found there were two types of offloading: (1) *self-referencing* and (2) *team-referencing*. Most of the participants flicked documents, images and entities from the private space of their own iPad to the shared screen space, but did not immediately use them in their thought processes. Instead, the participants merely used the spatial affordance of the tabletop to store information for later exploration or to bookmark potentially important documents. Many participants mentioned that they employed the tabletop only for self-referencing. For example, participants often transferred documents to the tabletop when the documents included keywords or entities that were hard to remember, such as exotic names and phone numbers, in order to reference them later when they came across the entities in different documents. Interestingly, all participants used this approach to record important information instead of using the bookmark feature in the Foraging tool. On average, participants bookmarked only 1.8 documents ($\sigma=2.31$, median=1).

Flicking documents for the purpose of offloading allowed for opportunistic collaboration. Even though participants flicked documents for individual use, the shared documents led to unexpected collaboration opportunities. For instance, during the discussion, a participant flicked a relevant document (for self-referencing) on a tabletop, and thereafter slid that same document directly to another participant who needed it during collaboration.

Of course, there were teams who frequently flicked documents for the purpose of active collaboration or “team-referencing.” In such teams, each team member was well acquainted with what other members were working on; if they found possible interesting information for another member while they were reading a document, they flicked the document onto the tabletop or another shared display. While this behavior directed their individual and collaborative investigations, it occurred at the cost of “polluting” the shared display workspace with multiple documents and entities. Our observations concerning the main use of the shared displays in multiple display environments as a form of external memory resonate with observations concerning sensemaking on single large displays [1].

Need for Larger Space. Another notable observation in favor of multiple displays is the support for on-demand increase in screen space as needed for analytic activities. While foraging for information contained on the iPad, participants often required a larger concept map or needed to open multiple documents simultaneously. On the iPad, such an application will usually take up the whole screen; this was perceived as beneficial to direct attention, focus and thinking [102]. However, the inability of the device to support viewing larger concept maps and multiple documents simultaneously was a key barrier to the use of the device for visualization or analysis-related purposes. One user commented:

“I could access only one document at a time with an iPad, but I often wanted to check more than two documents at the same time. Also, I needed to see relationships between entities across different documents but couldn’t read multiple documents on an iPad. In response, I spread multiple documents on the tabletop by moving them from my iPad.”

Participants could extend their workspace physically by flicking their content or entity from the personal tablet screen to the tabletop. This lightweight gesture interaction allowed participants to use nearby displays as extensions of their personal displays. No one attempted to reverse this gesture and flick information from the large display to an iPad.

Participants strongly agreed that VisPorter’s gestural interaction to move objects was extremely useful and allowed them to take advantage of nearby screens to transfer data and tasks; (4.6/5.0, $\sigma=0.67$, median=5). Almost all participants flicked the contents of their personal display onto a nearby larger screen in order to explore multiple documents or visualize them on a large display capable of displaying more detail than is possible on an iPad.

4.4.4 Objectification of Information

Objectification of information [102] occurs when users appropriate a physical object as a “carrier” of a specific thought or concept to be shared in a direct, transparent and quick manner. In such instances, users focus solely on the material being handled (e.g. the concept), as opposed to undertaking procedures to share information divorced from the

meaning of the object itself. In our case, objectification refers to how participants assigned meaning to devices. They associated concepts to particular devices, and used these “physical carriers” to expand their thinking.

We found that after organizing many documents that were related to a specific entity on a single display such as the iMac, this display was then regarded as a physical entity or representational proxy when team members discussed that topic. For instance, after collecting or moving all documents related to a suspicious person in the dataset onto the iMac, participants frequently pointed to and referred to the iMac as the suspicious person when discussing relationships among events involving the person. Three teams (G2, G6, and G8) displayed this interesting association. This type of physical referencing facilitates efficient communication among people [114]. In the interview session, one user commented on this facilitation.

“After collecting many related documents in iMac, I found that one guy was involved in several issues and events. Just calling him didn’t seem sufficient when we discussed him. I felt like that the large quantity of information related to that guy, and iMac becomes a physical icon. When I need to discuss something relevant to him, it seems easier and more natural to map or point to that iMac.”

4.4.5 User Feedback

In our post-session survey, VisPorter was very positively rated for finding hidden hypotheses in the dataset (Figure 4.11). The question “Rate your enjoyment when using the system?” rated an average rating of 4.0/5.0, with $\sigma = 0.85$. The question “How useful was the system in finding answers?” rated an average rating of 3.6/5.0, with $\sigma = 1.16$. On the other hand, for the question “How much did the system lengthen time required to analyze the data?” received an average rating of 1.9/5.0, with $\sigma = 0.79$.

In the interviews, the majority of the participants gave mostly positive feedback about the physicality and spatiality of VisPorter on multiple displays.

“I liked the idea of using my iPad to analyze each section of a document and then dragging it to the large display to organize information spatially.”

“The key advantage of this tool is that I am able to physically retrieve the information based on its place on the screen.”

“It was beneficial to be able to lay out data in multiple large displays. It also made working with a team faster, since we weren’t all looking in one place.”

Conversely, a few of participants felt stress using multiple displays due to the lack of information-management features across multiple displays.

“Many large displays are distracting and it is difficult to find specific information if too many documents are displayed.”

“I feel very insecure, because I was always afraid that the information on the screen would disappear. It’s easy to store information when you write it down. Then, when you want to retrieve the information, just get the paper. However, with multiple screens, we can’t easily record the information.”

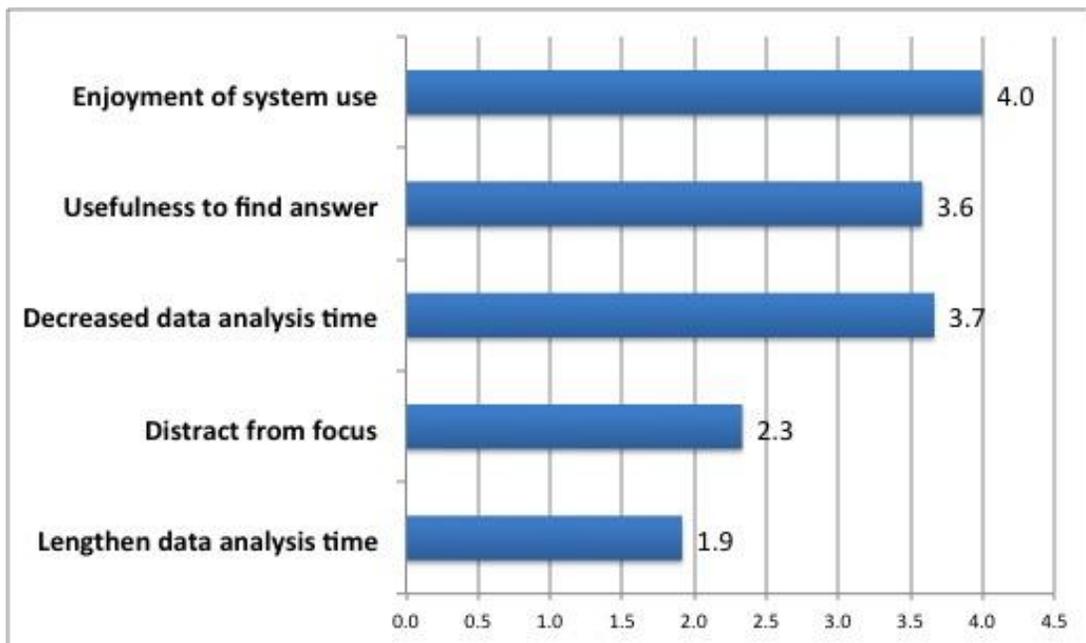


Figure 4.11. User feedback in the post-session survey (1-5 scale).

4.5 Discussion

4.5.1 Performance Factors

After the 1 to 1.5-hour analyses, six of the eight teams successfully discovered the overall situation, and seven teams successfully determined the key player in the dataset.

However, from the results of our study (see Table 4.1), we identified different collaboration styles and factors affecting the performance of the teams.

Specifically, we found G1 exhibited very low performance due to lack of information sharing and awareness of the other users' analyses. While G1 used FS (Fully Shared) and all of the team members shared only a single tabletop, they neither shared their findings actively on the shared display, nor tried to connect pieces of information different members had found. For instance, G1's members concentrated on individual analyses using a tabletop and each team member had different hypotheses than the other team members. As a result, G1 provided considerable misidentified information, yielding the lowest score among the teams.

Also, it is worth noting that the amount of exchanged (transferred) information between displays also appears to influence performance. The total number of documents transferred by each team ranged from 11 to 67, while the total number of entities ranged from 0 to 102 entities, which had an impact on analysis results. We observed that the group who shared and transferred more information across displays seemed to produce better results. In comparing groups that had the lowest and highest scores, we can see that the two high-scoring groups (G4, G8) exchanged a larger number of documents and entities between their individual devices and the shared large displays. They also employed more displays than the teams that received the lowest scores (G1, G3).

We also examined how objectification behaviors might affect their scores with multiple displays. However, the small sample size did not allow us to identify any significant correlation between the scores and this interesting behavior.

4.5.2 Deciding Better Analysis Strategies

In previous sections, we have shown various analysis tasks and patterns for visual analysis in display ecologies. We found that these analysis activities tend to be heavily dependent on user-specific strategies and intentions at the outset of the analysis. In this sub-section, we discuss how and why users decided to employ the initial analysis strategies they selected, what factors influenced their decisions for certain strategies, and what user analysis strategies appear to be the most effective.

We observed that users conducted a series of processes to decide on the analysis strategies they eventually employed. There were five decision patterns that participants used to determine their analysis strategies.

- *Negotiated:* Six teams (G2, G4, G5, G6, G7, and G8) used both preliminary analysis to familiarize themselves with the dataset, and then employed a negotiation phase to determine the optimal analysis strategies among collaborating users at the beginning of the analysis sessions. In the negotiation phase, each team determined the initial coordination strategies for analysis. Specifically, the members of these teams first read the documents on the tablet individually, and after all team members were somewhat acquainted with the dataset, the teams began discussing some preliminary findings. They then decided how to work together or coordinate the data across the different members or displays. Negotiating the division of tasks was generally conducted through face-to-face interactions based on user interests.
- *Emergent:* Among the negotiated teams, two teams (G2, G5) spent longer in the preliminary analysis phase. In these those two teams, some users began to spatially organize the documents over multiple displays without discussion, which prompted other participants to follow suit. We observed that the preliminary analysis allowed users to recognize the need for specific cohesive strategies to avoid redundancy and reduce the complexity of the analysis.
- *Leader-driven:* Throughout the analysis sessions, two teams (G4, G6) clearly had a team member who played the role of the leader. We observed that members of those teams were more successful in monitoring the work of their colleagues, as evidenced

by the fact that they steered each other towards more productive lines of investigation. These users frequently moved between displays and checked each other's progress and findings.

- *Evolve*: We observed that G2, G4, G6, G7, and G8 changed their collaborative analysis strategies as the analysis progressed. They later changed their strategies based on their evolving needs as more information became known and understood. For example, G7 changed their analysis strategies from SI (Strictly Individualized) to SD (Semi-Divided). At the post-study interview, a team member reported that the main reason for modifying his collaboration style was due to the fact that being aware of his colleague's findings and organizations was useful for his own line of investigation. Therefore, the team members frequently checked each other user's findings; moreover, they ended up sharing a single display rather than working on their own devices separately.
- *None*: It should be noted that two teams (G1 and G3) did not engage in face-to-face negotiations. In fact, participants in these teams shared very little verbal communication with each other, but instead concentrated on individual analysis on only one shared large display without using other large displays throughout the session. Due to lack of collaboration and awareness of other users' analyses, these teams took more time to form shared insights and conclusions. The fact that they were slower to synthesize and merge individual findings at the end of the analysis session led to the lower performance of these teams.

Because this study was basically exploratory in nature, participants were not given any tutorials on recommended analysis strategies with multiple displays. However, we confirmed that the multiple discretized spaces afforded by the display ecologies provided a natural way for them to spatially organize different information and tasks across displays. In fact, we observed that five of the eight total teams involved in this evaluation spontaneously took advantage of spatially organizing information items without any prompting or guidance on how best to carry out the analysis in display ecologies.

Performance levels were linked to user strategies in terms of how they collaborated in creating a “space to think” with multiple displays. In other words, it was crucial for the participants in this investigation to collaboratively create semantic structures and actively use the multi-display spaces as external memory. With respect to the creation of semantic structure, most of the high performance teams (“Negotiated” teams except G4) collaboratively formed semantic structures by spatially organizing documents across display. In terms of creating external memory, the majority of team members also actively flicked more documents and concept maps onto a shared screen for bookmarking and sharing information with other members. In contrast, some lower-performing teams (G1, G3) opted for “None” strategies, in that they confined their analyses within a single shared display without actively organizing their findings across displays. As such, the users in these teams did not collaborate in creating more meaningful structures across multiple displays.

In short, one of the biggest advantages of collaboratively creating a “space-to-think” via the use of a display ecology is the significantly enhanced awareness of other users’ analyses and findings. Both tasks (creating a semantic structure and using display spaces as external memory) rely on users integrating and combining different findings for better awareness and, ultimately, for moving an analysis forward. This investigation confirmed that the use of a display ecology enhances users’ intuitive collaboration activities and improves opportunities for finding common ground during an analysis.

4.5.3 Spatial and Physical Actions

VisPorter was designed to enable people to distribute knowledge and ideas around the physical space. Spatial organization of collected information on displays was very fluid on VisPorter with multiple displays (D1). Also, the lightweight gesture-based techniques used to move objects between devices supported by D2 made it possible for users to perform all of the cross-device activities observed in the study. Throughout analysis sessions with VisPorter, participants used physical navigation extensively to forage documents on the displays (Figure 4.12). For instance, participants frequently re-found documents by physically navigating the multiple display space. A participant observed

that the experience of foraging documents in VisPorter was very similar to finding information from piles of papers on different desks. In many cases, users did not even use the keyword search feature, but instead tried to find items through physical navigation. During the post-session interview, users commented that because documents were spatially organized across the displays, they could rapidly pinpoint the spatial location of the documents on the different screens. One participant stated:

“I could not remember how to spell specific keywords when attempting to re-find documents, but I could remember where the information had been placed.”



Figure 4.12. Cross-device referencing with physical navigation. The user in G4 analyzed the concept map on his iPad and text documents on the tabletop. He used physical navigation to scan the documents on the tabletop rather than use the search feature.

4.5.4 Opportunistic Activities

VisPorter extends the analytic workspace opportunistically, enabling additional externalization and organization of information as necessary. Opportunistic activities were enabled because the participants did not need to focus on memorizing the data—instead only flicking and organizing it (D2 and D3). They naturally offloaded information using the tools at hand. We observed that the appropriation of personal and shared spaces was improvised according to the participants’ needs. As evidenced by offloading information activities to large displays based on user needs and preferences, the role of each display and the user’s activities continually underwent transformations among different displays during the analysis sessions as needed. As mentioned, the tabletop was generally recognized as a public space, but participants also used it as an extension of their personal displays to see multiple documents and large concept maps.

4.5.5 Promoting the Objectification of Information

Many current collaborative sensemaking tools based on single displays (e.g., [8]) embody a model of collaborative sensemaking where users perform collaborative work with a shared focus and simultaneous individual control of visualizations on separate single displays. In these tools, the collaborative sensemaking is mostly restricted to the single shared virtual space. Conversely, VisPorter allows users to collaborate using interconnected devices that separate individual and shared work with natural physical affordances. This characteristic of VisPorter promotes the objectification of information, which enables users to regard concepts through physical devices as efficient representational proxies. In essence, the device becomes the information. In this instance, objectifying all the information related to the suspicious character as a physical display allowed them to consolidate all of the attributes of that character as a single unit—and then physically reference that unit while deliberating the character’s role on the plot.

This form of objectification is distinct from the notions of *object-orientation* [70] in that the object represented is conceptual in nature (e.g., the suspiciousness of the person) and the representation itself is a physical device, not just a visual representation on a display.

4.6 Summary

In this chapter, we presented VisPorter, a visual analysis tool with intuitive gesture interaction for information sharing and spatial organization in a display ecology. It strives to deliver a seamless experience for collaborative sensemaking across varied devices. The system embodies the idea that the multiple devices should operate as an ecology of mobile devices, desktops, and large displays for organizing and analyzing information. In this ecology, each device is afforded different analysis tasks (e.g., personal displays for foraging and large displays for synthesis), and has different effects on how participants make sense of information. We proposed a set of design principles derived from prior studies of single and multiple display systems. Our study of VisPorter with participant

teams, based on these design principles, showed that the concepts of “space to think” [1] extend usefully to display ecologies that support:

- **Flexible work division:** VisPorter supports flexible work division approaches by allowing team members to coordinate different analytical tasks among physically separated displays.
- **Cross-display semantic structures:** VisPorter allows team members to organize documents and concept maps onto different displays, based on the device capabilities and visualization needs as well as different entity types.
- **Extension of display space:** VisPorter enables users to move all information objects including text documents, images, and concept maps throughout displays in the workspace by lightweight gesture interactions. These approaches allow users to extend their workspace as necessary by transferring individual information or concept maps from the personal tablet to nearby available large displays.
- **Objectification of information:** VisPorter presents the greater opportunity for “objectifying” information using the physicality and spatiality that the display ecology affords.

Based on our analysis of participants’ use of VisPorter, we validated a set of design principles for multi-device systems that appeared to provide a cohesive and integrated experience. The results of our study inform the design of new sensemaking tools to help people leverage space in display ecology scenarios. Our future research goal is to improve the robustness and usability of the system, and to study the effects of using such a system empirically with a greater longitudinal basis.

5 A Comparison of Two Display Models for Collaborative Sensemaking

The current proliferation of mobile devices and large high-resolution displays offers new opportunities for both personal and collaborative sensemaking. If multiple displays and devices could function in a unified manner, would the sensemaking process be distributed in such a way as to generate cognitive (and other) advantages? How would such a “distributed” model compare to the current model where collaborative sensemaking occurs within the boundaries of a single display?

Prior literature has highlighted several benefits associated with the use of multiple displays and devices for data analysis and sensemaking —principally due the variety of affordances inherent in a display ecology. For instance, the multiplicity of devices exploits the human capacity to use spatiality and physicality to make sense of information [6].

The separate and common discrete spaces of the various devices also facilitate the division of tasks across different displays and among team members.

In contrast to this multi-display environment, the customized format is for groups to engage in sensemaking within the confines of individual computers with shared focus and simultaneous control of information. While this represents a tremendous improvement over past models of users working on isolated devices that do not have access to common shared information, we believe that there are greater benefits to be gained from allowing sensemaking to occur within an ecology of display and devices. In this chapter, we investigate the benefits that users may derive for the process of sensemaking to allow users to distribute cognitive resources across physical space. To this end, we compare the use of two systems, *VizCept* [8], which will be described in detail later, and *VisPorter* (Chapter 4). These systems both support the above-mentioned collaborative sensemaking environments with multiple displays, although in different ways.

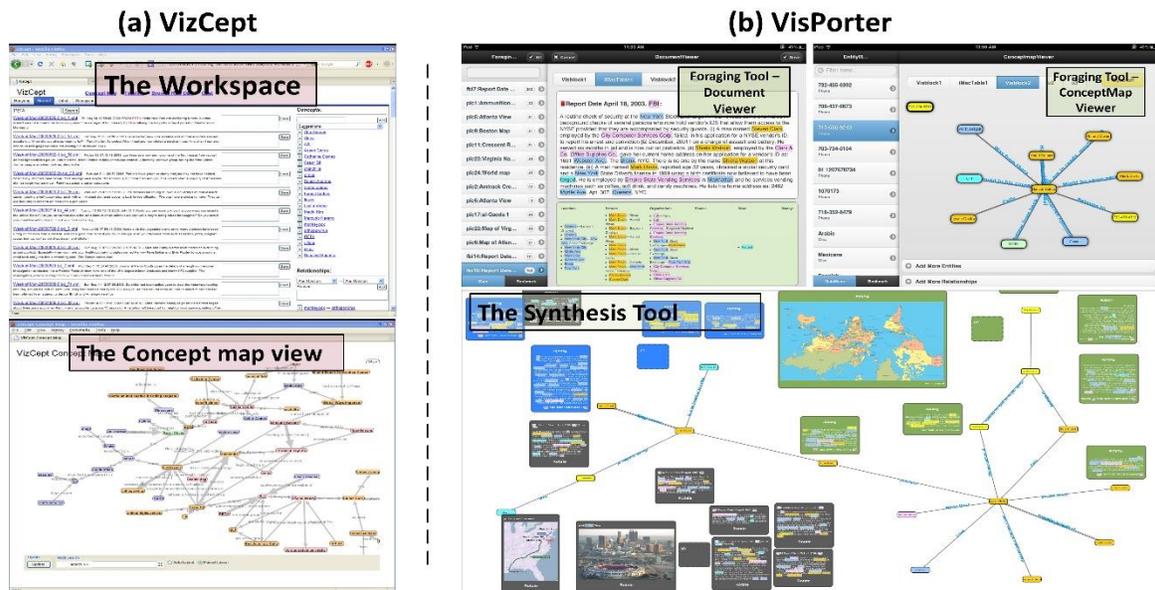


Figure 5.1. The two collaborative sensemaking systems used in our comparative studies.

5.1 Two Display Models

In this section, we describe the design of our prototype multi-display visual analytics systems, VizCept and VisPorter, which are two contrasting models for shared visualization on single displays and unified multiple devices, respectively. The two systems are based on a common framework that we will explain first prior to describing the particulars of each system.

VizCept and VisPorter are visual analytics systems designed to support co-located collaborative analysis of textual data by providing shared focus of information through concept maps. Both tools emphasize seamless transition between individual and collaborative analysis, which is an important foundational concept for group work [49]. Both VizCept and VisPorter consist of two types of sensemaking tools: the foraging tools and the synthesis tools. Each of these is primarily designed to support different stages of the sensemaking process. These two tools are directly analogous to the two loops in the model of the sensemaking process [79]: the Foraging and Sensemaking Loops. It has been shown that the division of the sensemaking process into these two loops can be beneficial for collaborative sensemaking, but that the two loops are highly interconnected [18]. Both systems include the following common features:

Foraging tools. The *Workspace* of VizCept (Figure 5.1a) and the *Foraging tool* of VisPorter (consisting of the *Document viewer* and the *ConceptMap viewer* in Figure 5.1b) are the main components for data exploration, providing keyword searching and document content browsing. In VisPorter, each document is automatically parsed for entities using the *LingPipe* library [105] (Figure 5.1b upper left). The Foraging tool and the Workspace also allow the user to specify the relationship between the entities.

Synthesis tools. The *Concept map view* (Figure 5.1a bottom) of VizCept and the *Synthesis tool* (Figure 5.1b bottom) of VisPorter enable the visualization of global concepts and relationships that collaborating users have discovered. This visualization is shared among all team members. Nodes in the visualization represent concepts or

entities, while relationships among concepts are represented as directed edges with descriptive labels. The colors of nodes represent different users or types of entities.

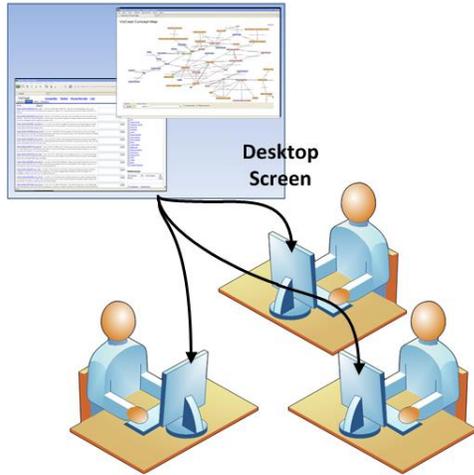
5.1.1 VizCept: Shared Visualization Spaces

VizCept [8] is designed such that each user employs individual devices such as laptops, tablets, or personal large displays. VizCept allows multiple users to distribute and parallelize analysis tasks on individual displays by foraging and collecting information individually. In this way, collaborating users can share and construct visualization through shared workspaces on individual displays (Figure 5.2a). Simultaneously, each user contributes to creating a shared concept map, which facilitates not only a heightened awareness of other users' progress, but also enhances the connections between individual findings and the collective work of the group. It must be stressed, however, that this system does not allow for any direct cross-device interaction. An analogy can be drawn with the popular *GoogleDocs* model, whereby each user accesses the shared document(s) on her own device, while being able to see updates by others in real time. The specific characteristics of VizCept are described below:

Interaction: The user interacts with the system through a conventional tethered interface such as a keyboard and mouse on the user's personal computer.

Concept mapping: Each user contributes to creating a global concept map, which facilitates an enhanced understanding of a given analysis task. The concept map helps to track valuable information in a one-screen view. Navigation strategies such as pan and zoom, or the manual/automatic layout (force-directed) of the concept map, can be applied individually on a shared concept map. The shared concept map provides awareness of the progress of the other users and the connection between one user's individual work and the work of the rest of the group (Figure 5.1a bottom).

(a) VizCept



(b) VisPorter

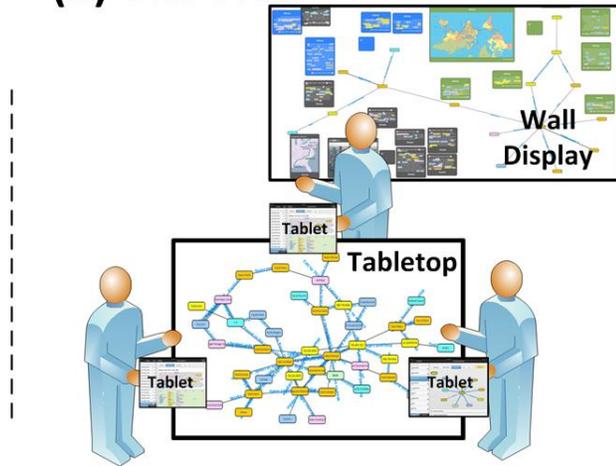


Figure 5.2. Two display models for collaborative sensemaking.

Scalability: VizCept’s multiple coordinated views (separate documents, pictures, concept maps, etc.) help users see different aspects of the same dataset on a single screen. However, multiple views and visual scalability of concept maps remain limited to a single display and computer.

5.1.2 VisPorter: Display Ecology

VisPorter is designed to facilitate information sharing between multiple displays using the physical reference of the displays. For instance, to transfer information from one device to another, users refer to the physical position of the target display. One can thus spatially distribute entities, concept maps, and documents across different displays—and then organize and investigate them further on individual or shared displays (Figure 5.2b). In VisPorter, information can be individually analyzed on one device and also shared with other collaborators and devices. Chapter 4 provides detailed characteristics of VisPorter.

Display: While all tools run on a single individual device/display in VizCept, VisPorter, in contrast, enables analysis across multiple individual displays (e.g., smart phones, tablets, laptops, etc.) and shared displays (e.g., tablespots, powerwall, etc.). However, the Foraging tool (Figure 5.1b top) is adapted to run on personal devices instead of on the shared devices. Conversely, the Synthesis tool allows users to take better advantage of large screens by organizing documents and concept maps spatially on the screen, and by enabling the integration of information and visualization items from multiple users and devices (Figure 5.1b bottom).

Table 5.1. Design characteristics of the two systems.

Design Characteristic	VIZCEPT	VISPORTER
Display	Individual devices at a time (laptops, tablets, personal large displays)	Multiple personal devices and shared displays (tabletops + wall displays)
Interaction	Desktop-bound mouse and keyboard, enabling only virtual navigation (zooming/panning)	Enabling spatial, gestural and physical navigation through touch-based interaction
Visualization (Concept maps)	Shared information, individualized concept map layout	Individual information + Individual layout; Shared information + Shared layout
Information Sharing	Automatic online updates	Online controlled/manual updates + Direct transfer of information and concept maps
Scalability	Limited to one screen/device	Users flexibly extend a screen space with other nearby screens
Awareness/ Progress Indicator	Through shared concept maps	Through visual scanning of associated displays and other users' actions

Interaction: See Section 4.2

Concept mapping: Users can create personal concept maps on the ConceptMap viewer (Figure 5.1b upper right), and then later merge them with the larger concept map on the Synthesis tool (Figure 5.1b bottom). The tap-hold gesture is used to transfer an entity or concept map across the devices. Therefore, multiple users can construct sub-concept

maps independently on their personal devices (*iPad*) and then combine them with the Synthesis tool of the shared display (e.g., wall or tabletop displays).

Scalability: In VisPorter, users are able to extend the device they are currently using with other displays and devices by moving objects (e.g., document, concept map, etc.) on one of their devices to another. The key characteristic differences between VizCept and VisPorter as described above are summarized in Table 5.1.

5.2 Study Description

As information sources and needs expand almost exponentially, data analysis and sensemaking are no longer processes restricted to formal domains like intelligence analysis. Accordingly, we designed and carried out two exploratory studies for the two different configurations (VizCept and VisPorter) of collaborative sensemaking described in the previous section, and compared qualitative results obtained regarding the sensemaking and analytic processes from the two studies. For both studies, we recruited 11 teams with three members each (total of 33 participants, 4 female and 29 male). All participants were graduate students in a variety of engineering disciplines at Virginia Tech. Although the participants were not intelligence analysts, they had basic knowledge of how to approach analytic problems from having taken graduate level Information Visualization classes; thus, they were familiar with data analysis procedures. Prior user studies in collaborative visual analytics have also made use of participants not formally trained as data analysts [41], [18]. The study tasks were common enough that they did not require any specialized knowledge.

Three teams were assigned to the VizCept study, wherein participants were asked to use individual devices (iPads or laptops) collaboratively (each participant had only one device). The remaining eight teams were assigned to the VisPorter study, wherein each participant had one iPad, but could access the following shared displays at all times during the analysis: a touch-enabled iMac, a tabletop and a wall display. Both the tabletop and wall display were made of nine tiled back-projection displays arranged as a

large 4ft by 6ft (3840x2160, 82.5 inch diagonal) horizontal or vertical surface screen with a PQ Labs' 32-points multi-touch overlay.

An intelligence dataset was used as a sensemaking task. Specifically, participant teams were asked to identify terrorist plots hidden within the dataset, which consisted of 41 documents. Additionally, we added 25 pictures to the original dataset related to important entities. All the team members were asked to come to the laboratory to perform collocated, synchronous analysis of the dataset, which lasted from 1.5 to 2 hours depending on the team. After completing the demographics questionnaire and tutorial session, each team then began the analysis task using the devices allocated in their assigned condition. The participants were asked to complete an answer sheet for a hypothesis solution with supporting evidence, including details (who, what, where, and when). After the analysis session, each participant was given a post-questionnaire to complete regarding their experience with the system, with particular focus on the analytic workflow they used to arrive at their solution. A group interview was subsequently conducted with all team members.

All analysis sessions were recorded (video and audio). Observation notes were taken by a researcher who remained in the experiment room. Screen activity was recorded for all work carried out using the Synthesis tool of VisPorter on the wall, tabletop and *iMac* displays; screenshots were taken at 30-second intervals. Additionally, all interview results and conversations during the collaborative analysis sessions were audio recorded and transcribed by the authors.

We employed a mixture of the multiple types of data for the analysis. We first consolidated our observation notes, interview transcripts, and post-session questionnaire results collected from the two studies by discussing and collating them on a whiteboard. Based on this process, we generated a set of key insights regarding the sensemaking process. We then conducted a validation procedure of those key insights by revisiting all other types of relevant data, including video and audio recordings of the sessions to document those activities in each display usage model. For instance, we identified users'

individual preferences for physical navigation in information foraging across multiple displays as a key insight, and then analyzed multiple teams' actual behaviors for physical navigation in recorded videos. If data from the various sources supported one of our key insights, the insight was considered to be supported. If data did not support a key insight, that insight was contradicted or not supported. We present in this chapter the key insights that remained as supported after all data sources had been analyzed.

5.3 Findings

The analysis workflow of VizCept was partially identified in [8]. Based on post-session questionnaires, observation notes, and interviews, we found that team members in both the VizCept and VisPorter conditions generally used a common analytic workflow consisting of five stages (Figure 5.3). Nevertheless, there were interesting differences within each of the five stages of the process. We present the key insights of the differences between the two study conditions within each stage below:

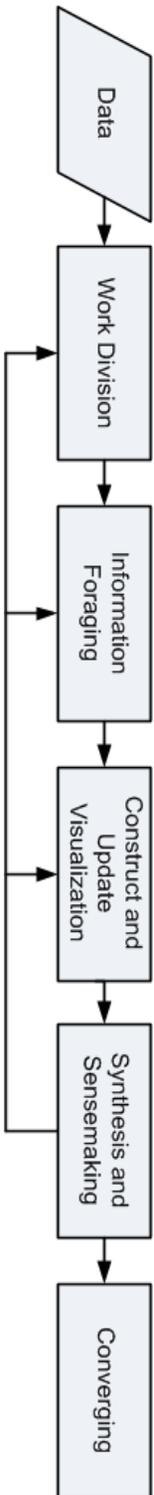


Figure 5.3. The common analysis workflow.

Stage 1: Work and Data Division.

This first stage generally consisted of the team working together and coordinating the data across the different members or displays. We observed a notable difference in the division of work between the uses of the two systems in relation to the process of how it was achieved.

Using VizCept, the teams relied only on communication methods that were external to the system in order to reach a consensus of how to divide work and data among the team members. Specially, they used verbal communication (i.e. oral negotiation, discussions of which information object belongs to which group) and textual means (i.e. Internet chats and annotations) to achieve their goals. In contrast, VisPorter generated a physical way of dividing work and data. The physical spaces of the different displays were used to divide and organize information. For example, pieces of data were assigned to specific screen spaces of the different displays. In addition to the division of data between displays, team members assigned themselves to different displays. For instance, in one team, the tabletop, wall and *iMac* displays were divided among each of the three team members and were used as individual workspaces in addition to the individual iPads. After discussion, the team then assigned the categories of data to a suitable large display based on content and entity type for further analysis by the associated member.

Stage 2: Individual and Collaborative Information Foraging.

Information foraging required participants to search for keywords in the documents, read them and identify new concepts and relationships. Under both study conditions, participants first read the documents loaded into the system individually, and after all team members were somewhat acquainted with the dataset, the team began discussing documents related to specific topics.

In the case of individual foraging, the two compared conditions differed in terms of how participants marked relevant information objects for later use during group analysis. With the VizCept condition, participants either added notes to important documents (i.e., as

annotations), or they used the bookmarking feature. In the VisPorter condition, however, participants used the space provided by the shared displays to support external memory of information for their own reference (self-referencing; Section 4.4.3). For example, participants transferred documents from their iPads to the tabletop when the document included entities that were hard to remember, such as exotic names and phone numbers, in order to reference them later when they came across these entities in different documents. Additionally, VisPorter users showed a clear preference for physical navigation when foraging for information or specific documents across the displays, rather than searching keywords on the individual iPads.

With respect to group information foraging, the two conditions differed in the way that insights about information objects were shared. Using VizCept, information sharing among the team members was initiated with oral file referencing to other team members, as in Stage 1. In other words, participants verbally referred to specific document IDs to ask another user to review the document. In contrast, when participants were using VisPorter, they merely flicked documents onto shared displays (Section 4.4.3).

Interestingly, this created other opportunities for chance collaborative moments. For instance, one participant flicked a document (for self-referencing) on the tabletop. He thereafter slid that same document directly to another participant's display during collaboration. Opportunistic interactions also occurred because participants did not need to focus on memorizing the document IDs, but only on the task of organizing the data and making sense of it.

Stage 3: Constructing and Updating Shared Visualizations.

Both systems use visual concept maps to represent associated thoughts from multiple users. In addition to the construction of their own concept maps, team members contributed to the buildup of the shared global concept map by creating, merging and refining entities and relationships based on their findings.

One key difference between VizCept and VisPorter is that the former shares all objects added to a global concept map indiscriminately. This synchronization (immediate

sharing) of concept maps had two contrasting effects. On the one hand, some members were hesitant to add premature results and concepts with the concern that they would further confuse the seemingly disconnected analyses, thereby hindering a productive line of investigation [8]. On the other hand, the synchronized concept map helped to create common ground among team members. For example, one participant found that no one else had added any concepts related to her concepts for more than 30 minutes. That prompted her to question whether her line of inquiry was wrong and to try and find the reasoning behind the information objects that the other members were adding.

In contrast to VizCept, VisPorter allows the participant to retain concept maps on individual devices locally, while concept maps on shared displays are global. In the VisPorter condition, we saw a more refined process of concept mapping—as evidenced by the fact that participants first narrowed their initial concept maps on their individual devices and only selectively flicked parts of their concept maps onto the global concept map on the shared display. Furthermore, in the VisPorter condition, it was evident that users were assigning concept map data to particular displays based on data type, display size and device capabilities (Section 4.4.2).

Stage 4: Synthesis and Sensemaking.

Synthesis involves participants integrating and combining multiple insights from all team members and developing a common series of insights. Based on the results from this stage, users decide whether they must return to one of the previous stages or proceed to the final hypothesis.

With VizCept, the formation of common insights was an additive process; in other words, insights from individual members were brought together through the concept maps and verbal communication (i.e., through oral explanations and internet chats). Verbal communication was the primary means to validate individual findings. In contrast, with VisPorter, synthesis occurred in a collaborative process that provided more awareness of others' activities and integrated cycles of common insight formation and

presentation. For example, the VisPorter teams worked together to organize document objects across displays or refine concept map nodes.

Stage 5: Converging.

In both models, convergence occurred when improvements to the concept maps were completed or when it was time to arrive at a common conclusion regarding the solution of the plot. In the VizCept condition, all of the teams allocated “presentation” time to each member to relate his or her conclusions/story. Each member explained his or her own cluster in the global concept map, as everyone gathered around one display. After all presentations, a common final story was agreed upon. In the VisPorter condition, this stage was brief as participants engaged in discussions to find common ground throughout the whole analytic process. So, the more formal “presentation” process played a less important role in the analysis. The physical space and engagement in spatial organization of documents/concept maps, afforded by multiple displays, changed analysts’ approach to convergence by tightly integrating the information synthesis and presentation stages.

5.4 Discussion

VizCept embodies a model of collaborative sensemaking whereby users perform joint work by having shared focus and simultaneous individual control of the dataset on personal displays. However, the collaborative sensemaking is mostly confined to the individual screen and verbal communication. Based on our investigations, VizCept did exhibit certain positive effects. For example, the system’s capabilities for merging collaborators’ thoughts/findings in a global concept map, as well as in single screens, facilitated monitoring the work of their partners. Conversely, VisPorter embodies a model whereby users collaborate using varied interconnected devices that separate individual and shared work with physical constraints. It enables people to distribute knowledge and ideas around the physical space where the displays take on meaning. The key differences in the use of the two systems for sensemaking that we elucidated from our study are summarized in Table 5.2.

Table 5.2. Key differences between the two models.

Stage	VIZCEPT	VISPORTER
Work division	External communication methods (speech, text chats)	Accountability for actions and physical assignment of data and members to specific shared displays
Individual information foraging	Annotations to documents/bookmarking	Flicking documents onto displays, Physical navigation
Group information foraging	Mostly individual with oral file referencing to team (e.g., document ID)	Opportunistic collaboration; flicked documents for self-referencing are also used for collaborations
Updating shared visualizations	Greater noise; hesitancy in sharing opportunities for common ground	More refined; selective sharing of only important information
Sensemaking and Synthesis	Additive effect of individual insights and verbal communications	More awareness of others' activities and integrated cycles of common insight formation and presentation
Converging	'Presentation' mode	

Two common characteristics of the sensemaking processes that occurred under the VisPorter condition are particularly interesting, which have to do with its greater emphasis on the “physicality” of that model. First, VisPorter facilitated immediacy in information sharing among collaborators, whereby users appropriated information objects to be shared and received in an immediate and transparent manner. In short, the focus of attention was on the material being handled. The second characteristic concerns a process that we call the “objectification” of information [102], which refers to how participants assigned meaning to devices (Section 4.4.4).

One of the most notable differences between the two models was how to share information with other collaborators. VizCept requires several indirect procedures in order to share information across displays—e.g., referring to specific document ID. In contrast, VisPorter enables users to share information immediately by flicking an information object from one screen to another. For example, if a VisPorter user wanted another user to read some documents, she simply flicked them on the shared displays,

instead of referencing to the document ID. In this way they were able to assign “thought objects” to particular devices, and used these “physical carriers” to expand their thinking. This concept is related to the idea of distributed cognition. For instance, after organizing related information on a particular display, the physical display device was regarded as a physical representational proxy for a collection of related data during discussions. In short, the device *became* the information. An illustration of this physical referencing is the frequent pointing gestures towards one display as the team members discussed a specific fictitious person in the plot (whose information were gathered on that display). In other words, objectifying all the information related to the fictitious character as a physical display allowed them to chunk all the attributes of the character as a single unit, and physically reference that unit, while deliberating the character’s role on the plot.

5.5 Summary

In this chapter, we investigated how the current paradigm of collaborative sensemaking differs from a prospective ecological model where all the displays in an environment develop roles and relationships for sensemaking tasks. The chief contribution of our work is to provide a qualitative comparison of two systems built for co-located collaborative sensemaking tasks that use different display and input arrangements. Although we found that the overall sensemaking process remained the same, we identified many differences employed within each stage of the process. A key benefit that the ecological model (VisPorter) brought about was in the greater opportunity for objectifying information afforded by the physicality and spatiality of the system. The differences between the two models as identified by this investigation can inform the design of new sensemaking tools or future groupware about how people leverage spaces in ubiquitous display/device scenarios. Our findings have not only significant implications for how future systems can be designed to motivate better collaborative sensemaking, but we also hope that it will generate discussion in the visual analytics community regarding the potential of new display ecologies and interaction approaches.

6 SAViL: Spatially-Aware Visual Links for Sensemaking in Display Ecologies

A typical sensemaking task requires an analyst to identify and understand various cognitive threads embedded throughout documents, images, and visualizations. As discussed in Chapters 4 and 5, when an analyst performs a sensemaking task with a display ecology, information of interest and analytical activities are typically scattered over different displays, thus requiring the user to switch intermittently among multiple foci of interest. The analyst must mentally connect and integrate diverse pieces of relevant information from different displays in order to generate a larger, coherent story. Thus, in contrast to a single large-display environment, the significant challenge associated with sensemaking using a display ecology is to maintain awareness of, and subsequently integrate, information from different data sources (often involving different visual

representations or data formats) over separate displays—several of which may be beyond the user’s immediate visual field [14].

This chapter focuses on visualization and interaction approaches to connect and direct a user’s orientation to important information located on different displays for sensemaking tasks. Many of the current sensemaking systems that employ multiple displays support information awareness and the ability to connect information on different displays via a strategy of synchronized highlighting utilizing brushing-and-linking approaches. For example, if a user selects specific keywords or visual elements on one display, associated keywords or elements can be highlighted on other displays with different colors or by enclosing them with boxes. Although these highlighted elements can make it easier for the user to distinguish important information from irrelevant data, there are two principal deficiencies:

First, the primary shortcoming of current highlighting approaches for multiple display systems is that the user must rely solely on memory to find various pieces of information which can become problematic when the amount of information and the number of devices are increased. If displays and workspaces are altered in a display ecology, analysts may forget the location of pertinent information. Furthermore, the highlighting approaches discriminate linked data items located on different displays with different colors or shapes, so users can perceive only a limited number of connections among these items on multiple displays [92], [93]. Second, the highlighting techniques are less effective for showing semantic relationships between multiple data elements scattered across more than two displays or data elements.

In contrast, visual links, which is a promising method for showing relationships between multiple pieces of information, has been widely investigated as a sensemaking tool—but only using a single display [26], [23]. The challenge, then, is how visual links might be extended in a multiple display environment, whereby analysts can be directed to important information across displays that are out of their immediate visual field. To address this important issue, we present Spatially Aware Visual Links (SAViL), which is

capable of elucidating the relationships among various entities, documents, images, or visualizations across different displays and devices (Figure 6.1 & Figure 6.2).

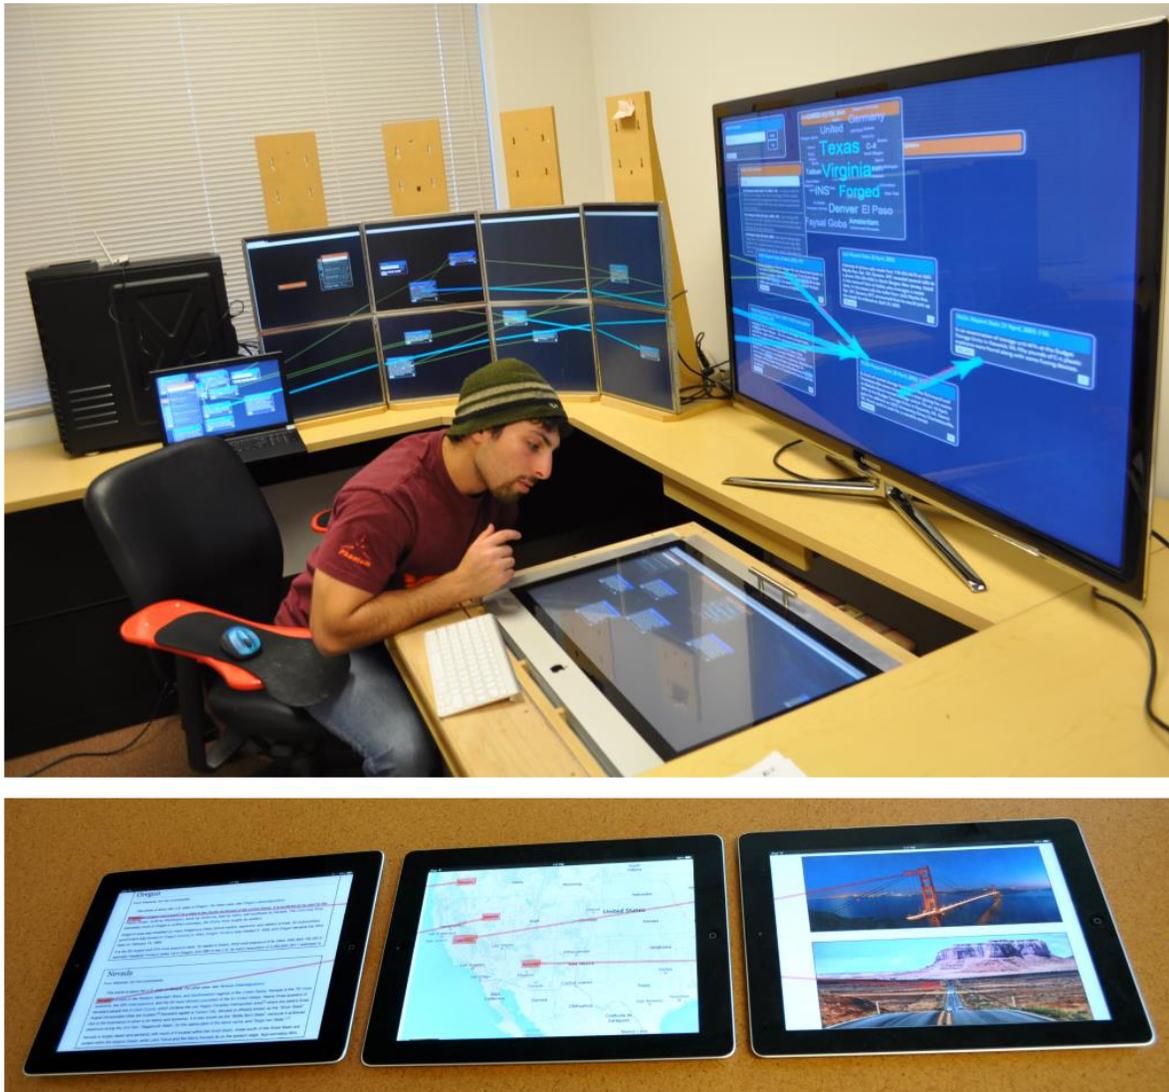


Figure 6.1. Spatially aware visual links for display ecologies.

This work contributes to the literature by describing a new visual link technique for display ecologies, which is expected to increase our understanding of the value of space for sensemaking with various displays. In particular, we expect to contribute to the field in the following ways:

Cross-device visual link techniques: The primary contribution of this work is to describe the design consideration and techniques for sensemaking, which utilize cross-

display visual links that help users connect and integrate scattered information across displays.

Impact on human sensemaking: For the second contribution, we extend prior investigations wherein users have employed single large high-resolution displays and their screen real estate for sensemaking [1] to the mixed-display environment. In the user study, we explore how cross-display visual links help users (1) become aware of information on different display and (2) recognize new connection between information across different displays. The results of this experiment will also show the impact and effectiveness of the cross-display visual links in a display ecology on the sensemaking process.

6.1 The SAViL Overview

SAViL was designed to provide simple visual links between diverse sources of information on multiple displays, creating spatially aware cues that may aid information synthesis. The design goal of SAViL was to construct an “integrated workspace” over multiple displays through cross-display visual link representations. Visual links can be drawn over displays to show relationships between information items located on different displays (i.e., explicit connection; see Section 3.3.3). SAViL was designed to facilitate an understanding of linked keywords and information items from multiple formats (e.g., text documents and images) across different displays. Specifically, it employs spatially-aware visual links to help analysts relate and locate information in display ecologies, as well as orient their attention to important information and the physical location of displays.

Each display in a display ecology maintains a separate, non-overlapping screen space, in which different information can be organized according to different attributes or data types (i.e., semantic substrates; see Section 3.3.1.2). Users can then spatially organize information across displays with the aid of the cross-display visual links. In this section, we provide a more detailed description of SAViL’s interface components.

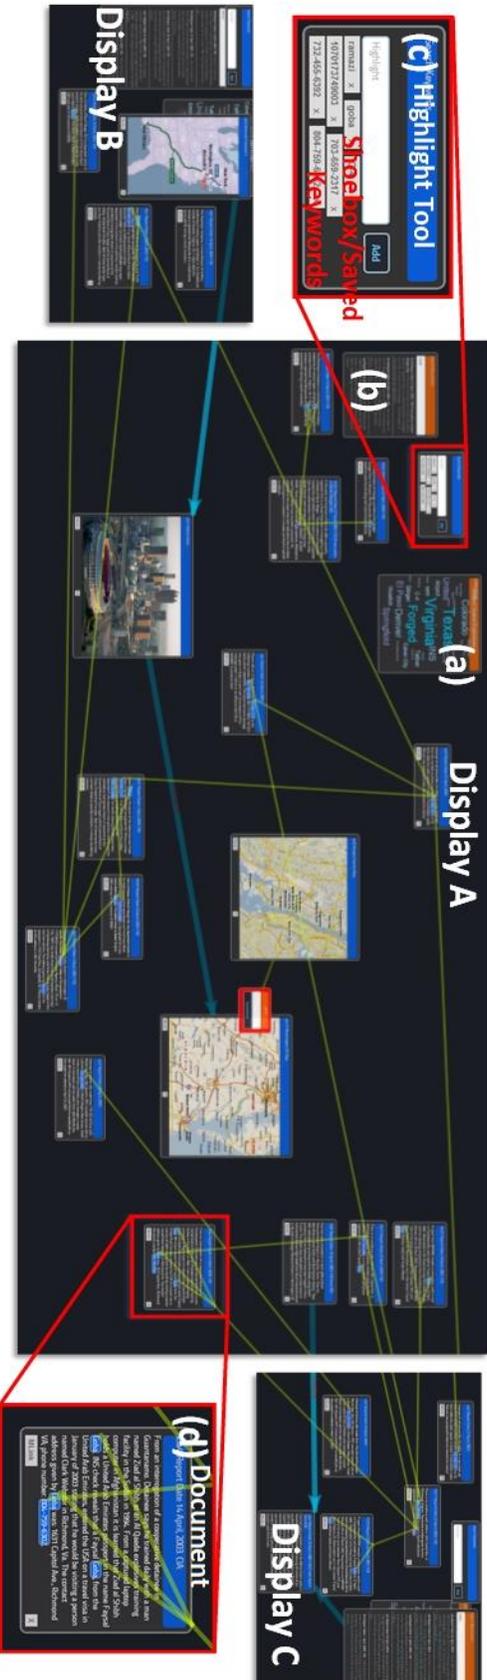


Figure 6.2. SAVIL with basic document analysis tools: (a) word cloud, (b) document search interface, (c) highlighting and shooboxing interface, and (d) document artifact.

6.1.1 Cross-Display Visual Links

SAViL's cross-display visual links are the straight lines from a source to multiple targets across displays and devices. These cross-display links employ the “partially out of the frame” approach advocated by Halo and Wedge [66], [94] (See Section 3.3.3.2). The theoretical foundation for this approach is based on the theory of amodal completion, which implies that a viewer will mentally complete the missing part of the link, even though only part of the link is visible [115]. Because these cross-display links are seamlessly drawn across displays (e.g., from a laptop to a wall display), they give the illusion of one continuous workspace utilizing different displays. The system supports both automatic and manual linking, as described below.

6.1.1.1 Automatic Linking

A link can represent a number of relationships between the source and target, depending on the level of abstraction or data type. If a keyword is selected, the selected keyword becomes the link source, which means that links are drawn to all target keywords on multiple documents across displays (Figure 6.2 & Figure 6.3). Based on personal preferences, a user can select from among four cross-display visual link approaches: (1) keyword/entity linking, (2) line bundles to document, (3) line bundles to display, and (4) document linking.

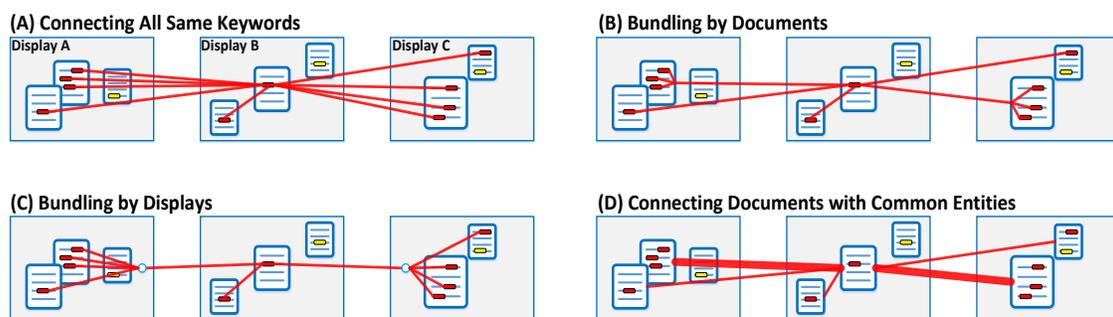


Figure 6.3. SAViL cross-display links. Each rounded box represents a document, and small red and yellow boxes represent entities. A user clicks an entity in a document on Display B and every same entity on different displays is automatically connected.

Keyword/Entity Link. The Keyword links are created automatically when users click on document keywords. The goal of this type of visual link is to keep users aware of how

keywords are related and spread over multiple displays (Figure 6.3a). This approach may help the analyst develop greater awareness of the number of entities of interest that occur in one or more scattered documents across displays. However, this type of link could introduce more visual clutter as the number of target keywords increases.

Line Bundles to Document. Using a hierarchical relationship system between entities and documents, SAViL can bundle multiple links from each document that contain the target entities [116], [92]. All of the internal targets within the document are then connected from the bundling point, which is an intersection point between the bounding box of the document and the link from the source (see Figure 6.3b). As we can see in this figure, this system reduces the number of links that bridge displays. Additionally, this approach facilitates the identification of keyword frequency among documents.

Line Bundles to Displays. This approach also employs hierarchical relationships among entities, documents, and displays. In other words, the link source is still a single entity, but the link target becomes any display that contains both the keywords and the documents. This can further reduce the connection lines across displays (Figure 6.3c).

Document Link. According to user preference, each document can show a connection line to other documents; this indicates how many extracted entities are shared between the link source and target documents. Varying edge thickness is based on the number of co-occurring entities between the two documents (Figure 6.3d).

6.1.1.2 Manual Linking (Annotated Links)

In addition to automatically linking between keywords, analysts can manually create relationships and annotations between two documents, even across multiple displays. For example, in our prototype system, the analyst first selects the source document and clicks the link button at the bottom of that document; this brings up a connection anchor icon, as seen in Figure 6.4a. To create relationships between two documents on different displays, the user simply drags the anchor across displays (Figure 6.4a) and places it on one or more target documents (Figure 6.4b). When the “connection” button is clicked on

the linking UI, the overlapped document becomes a link target, and the connection link and its label for the relationship is shown across displays immediately (Figure 6.4c).

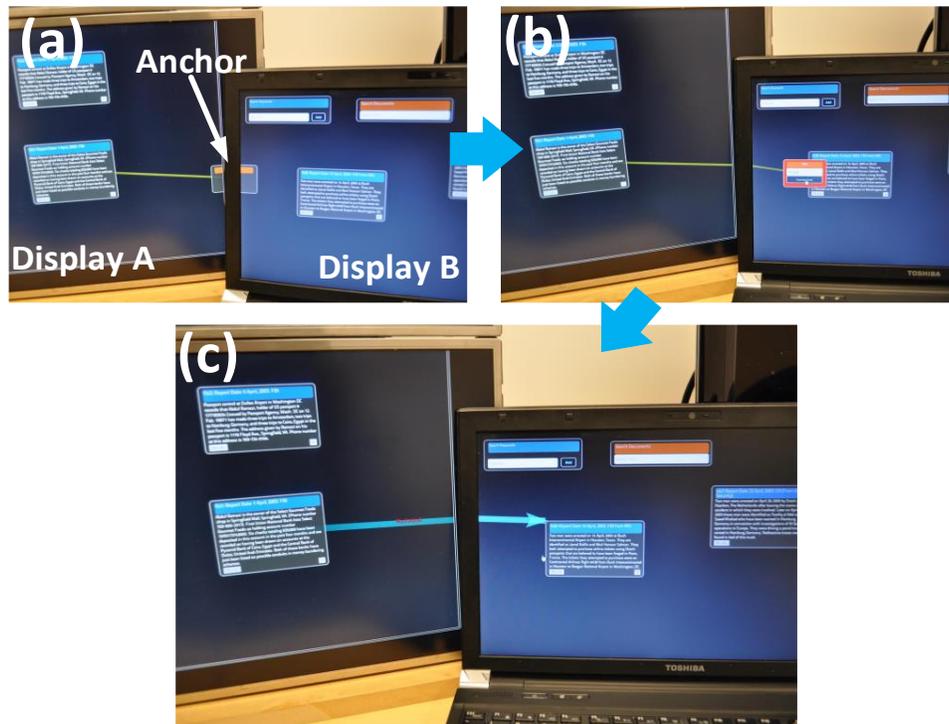


Figure 6.4. Manual linking. From left-to-right: (a) a user drags the anchor across two displays, (b) place it on a target document, and (c) a manual link is drawn across the displays.

6.1.1.3 Supporting Spatial Awareness

If a user drags a document to a different location around multiple displays, all of the links connected to that information artifact (e.g., documents and images) are updated across displays according to the new location. For instance, when an analyst changes the spatial layout of connected documents between two displays, all connected links are maintained and reoriented, regardless of where the analyst moves one or more documents. Also, in the case of portable screens, the system is capable of tracking the physical location of that device using the motion-tracking system and updating the links appropriately (Figure 6.5).

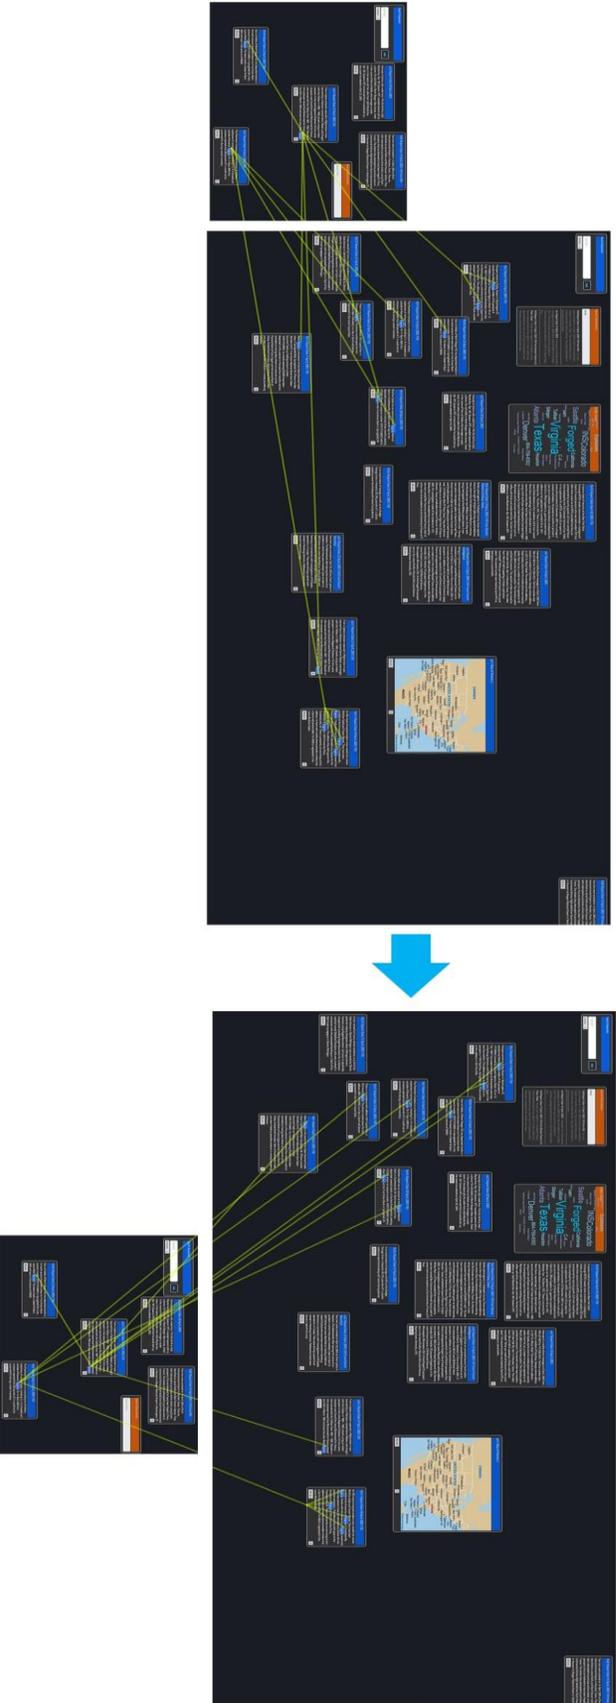


Figure 6.5. Support spatially aware links. A small display around a tabletop display is moved to a different location and the cross-display links keep following the location of the display.

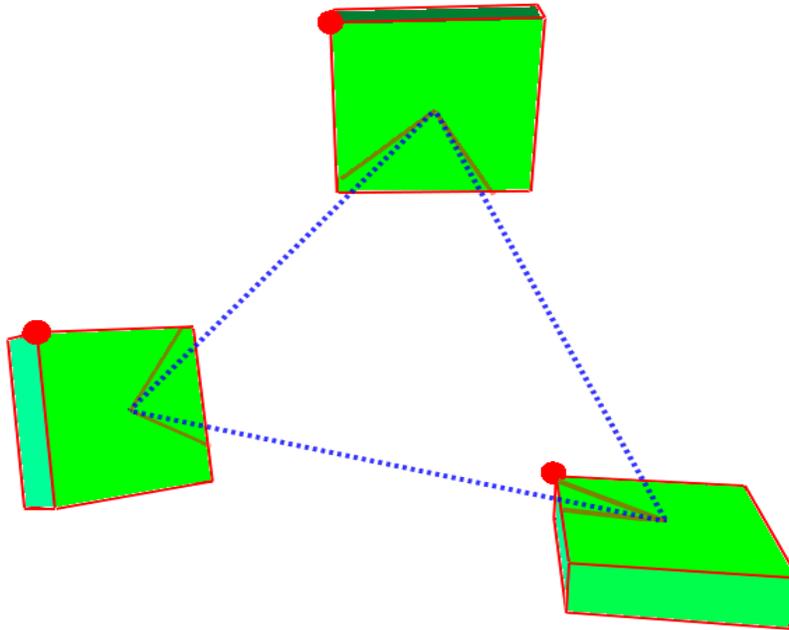


Figure 6.6. Drawing SAViL from 3D physical space to 2D screen. Green boxes represent the displays laid in the space, the red dots represent the position of each display in the space, blue dotted lines represent the 3d virtual lines, and solid red lines represent projected visual links on each display.

6.1.2 The SAViL Drawing Algorithm

The SAViL drawing algorithm performs the following functions (Figure 6.6):

- 1) Calculates the 3D physical position of documents, images and keywords based on each display's position (Figure 6.6 red dots) and rotation information
- 2) Calculates the 3D virtual links (Figure 6.6 blue dotted line) between documents, images and keywords
- 3) Projects the 3D virtual links into each display plane and calculates the 3D intersection points with each display's boundary
- 4) Calculates the relative positions between intersection points with the display top-left corner (Figure 6.6 red dots) and maps that information back into 2D position based on a display's size and resolution
- 5) Draws links in each display between the documents using boundary intersection points (Figure 6.6 solid red line)

In some scenarios, users can re-arrange their mobile displays in physical space in order to facilitate specific sensemaking tasks. In such cases, users can either manually determine the position of each display with the user interface (similar to the Screen Resolution applet in MS Windows), or use the motion-tracking system to track each display's physical position and rotation information automatically.

6.1.2.1 Size Adjustment

Due to the fact that displays differ in size and resolution, they are likely to have different pixel densities, which could be problematic for the consistency of the object sizes. Once the visual link is drawn across two different displays, visual properties such as a visual link's line width or the variable font size of documents will be adapted depending on the properties of the available display (e.g., pixel density). So regardless of the pixel density of each display, the size of visual links will remain uniform throughout.

We used PPI (pixels per inch) as the universal measurement standard for pixel density for various devices. In order to maintain the consistent visual link width $Link_i$ (inch) across different displays, we need to calculate the actual visual link width in pixels $Link_p$ in each display based on its pixel density.

Suppose, for example, that one display's physical diagonal is D_i inch, and its width and height resolution are W_p and H_p , respectively. Based on the Pythagorean Theorem, we calculate the diagonal resolution in pixels D_p :

$$D_p = \sqrt{W_p^2 + H_p^2}$$

Therefore, the PPI of this display P_p is:

$$P_p = \frac{D_p}{D_i}$$

So the actual visual link width in pixels $Link_p$ can be calculated as follows:

$$Link_p = Link_i * P_p = Link_i * \frac{\sqrt{W_p^2 + H_p^2}}{D_i}$$

6.1.3 Prototype System and Implementation

SAViL's basic role is to provide users with a sensemaking environment through explicit visual cues with which they could explore documents on different displays. To demonstrate the effectiveness of SAViL for document analysis, we created a basic web-based sensemaking tool that implements SAViL. The tool provides a suite of basic analysis tools to explore a large collection of text documents and pictures from a database. The primary interface for the SAViL prototype system is shown in Figure 6.2. Basic tools include a word cloud, document search tool, highlighter, and the shoebox tool to aid text analytics.

The SAViL technique is implemented with web-based client/server architecture. Specifically, the SAViL software infrastructure comprises multiple client applications that run on separate PCs or handheld devices in parallel. The server keeps these client applications synchronized through multiple managers to enable a coherent view for visual links and interaction across displays. In our infrastructure, the data between the client applications and server are exchanged in compressed Java Script Object Notation (JSON) format through Websocket or Ajax. In this section, we provide implementation details focusing on components of the SAViL architecture.

6.1.3.1 Client Applications

The client corresponds to a web browser on each device. The clients on different displays provide user interfaces and visualization views. The clients on multiple displays provide user interfaces and document artifacts from DOM (Document Object Model) elements and cross-display visual links are rendered by HTML5 CANVAS (Figure 6.7).

To connect keywords across displays through visual links on client applications, one must identify the position and size of each keyword across multiple screens. A unique DOM id is assigned to each keyword across all displays in the display ecology; thus, each keyword becomes a DOM element. The browser can provide and identify the position and size of each DOM element in a 2D screen coordinate. In the same manner, the system can

retrieve the position and bounding box information of each document artifact as a DOM element. However, because each screen maintains independent screen coordinates, the position information of the keywords and documents cannot be used directly to connect links across different displays. To connect links among elements on multiple screens, SAViL supports the World coordinate (3D), which represents our physical space in which displays and document artifacts are actually positioned. In other words, the world coordinate is the coordinate system of various visual objects and visual links on SAViL in the common 3D coordinate. The view of the client application on each display plays the role of a viewport to the world view, and all documents and highlighted keywords are generally managed in the world coordinate. For example, to draw and show visual links that connect objects between two displays, the coordinates of the visual links are converted from the world coordinate system to the 2D screen coordinates on each display (Section 6.1.2). Each display also checks its bounding box to determine if the objects and visual links are within their viewport. If an object is within the viewport, the client application converts the world coordinate of the object into its screen coordinate and display it. Our client applications are implemented in JavaScript, HTML5, CSS3, and the DOM-based linking approaches are generalizable to any type of web page, thereby enabling users to connect on any webpages.

6.1.3.2 Synchronization Server

The main role of the server is to synchronize and broadcast document and application state to client applications on different displays. It also keeps track of viewports on each display, as well as the location and configuration of the different client displays. Specifically, the server consists of three distinct managers.

Device Membership Manager: Our system allows users to organize the workspace dynamically with multiple displays. If an analyst starts a SAViL client application, the analyst's display is automatically included in the ecology, and the device manager reports the addition to the remaining displays. The membership manager allows a user to dynamically add or remove devices during execution (i.e., one device enters or leaves the analytical environment), while still being aware of device membership in real-time. Thus,

when a display is added or removed, the connected links across displays are immediately updated by reflecting the new display and its contained documents. The manager assigns a unique display id to each display and keeps a list of display members in the ecology, thereby enabling an analyst to reconfigure the workspace physically as needed, depending on the task at hand.

Artifact Manager: The artifact manager keeps tracks of the status and location of document artifacts—e.g., moving, removing, creating, selecting, etc. As mentioned, each artifact, such as a highlighted entity or document, has a unique DOM element id and this id is used to identify and report the changing status of specific artifacts across multiple displays.

View Manager: The view manager is responsible for spatial co-awareness between multiple displays. The view manager identifies and decides each display’s physical location. If a user should change the physical location of that device, the viewport of the display is also changed and then propagated to all other display views and visual links. The view manager retrieves the physical position of each display from a motion-tracking server or a manual layout UI based on the actual position of displays.

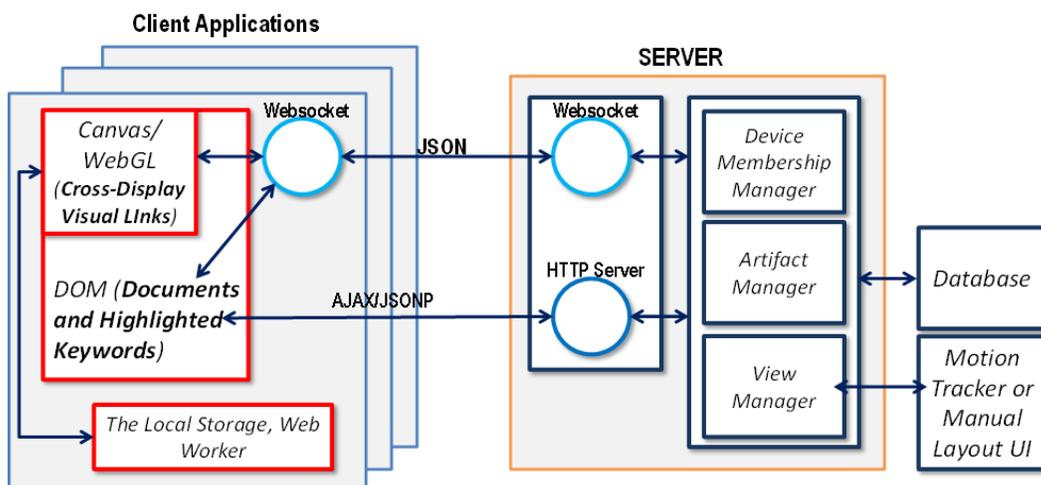


Figure 6.7. SAViL client/server architecture.

6.1.4 Usage Scenarios

We describe a sensemaking scenario for the investigation of wildlife law enforcement personnel and endangered species issues. The actual dataset available to our hypothetical investigators is a visual analytics dataset [103], which includes approximately 1700 files encompassing intelligence reports, news articles and pictures. The following fictional scenario is provided to illustrate the potential of SAViL.

Noah is a government employee who investigates illegal possession of endangered animals. In order to synthesize a significant amount of diverse data, Noah decided to utilize a display ecology consisting of one large display wall, a tabletop display, and a laptop. He initiated the analysis by searching keywords of endangered species, which enabled him to locate and open many relevant documents. Using visual links, Noah then grouped frequently-appearing terms or topics (e.g., persons of interest and location of suspected crime) and their parent documents. By simply clicking entities on the documents, Noah created visual links that then connected entities of interest across displays. He was also able to judge the importance of documents by identifying multiple links from the co-occurrence of different entities in the document.

The multiple-display configuration allowed Noah to organize the data based on device-specific capabilities and visualization needs. Thus, after quickly perusing many documents, he distributed and organized analysis tasks and data on his three available displays based on (1) people, (2) locations and events, and (3) organizations. This approach facilitated the distribution of analysis tasks across different displays so that he was able to work independently on different issues with different displays.

As the clusters across different displays increased, Noah sought to better understand the relationships of the various documents and clusters he had identified. Using the wall display, he first determined how people might be related to each other and to his search of interest. In so doing, he noticed a recurring name in many of the newspaper articles on his wall display. This person of interest, “rBear,” happened to be a famous pop star who openly espoused conservation wildlife issues. By clicking rBear’s name in a document,

Noah was able to link to other information related to him across different displays. Noah then simply followed connected links across displays to find other relevant documents—for example, related to rBear’s property ownership. In fact, utilizing his tabletop display, Noah was able to identify several co-occurrences of entities related to endangered animals and an animal sanctuary located north of San Diego on the tabletop display from the rBear documents he had previously placed on the wall display. Based on the linked information, Noah found a document suggesting that rBear was actually the “behind-the-scenes” owner of a big animal ranch. This seemingly contradictory information made Noah suspicious about the man who by many accounts championed wildlife protection. Because of rBear’s association with an exotic animal facility, Noah grew increasingly suspicious of rBear and decided to further investigate whether rBear was smuggling and reselling endangered animals.

During his investigation Noah frequently switched to different displays to organize and read relevant information. However, if the display had been altered in any way, he would have a natural tendency to forget about pertinent information on another display. To counteract this possibility, when Noah believed that certain documents might be related, he immediately made connections between documents located on different displays using annotated (manual) links. For example, Noah created some annotated links to a document showing that “rBear” had an alias, “Bert,” which he had found using another display. These annotated links helped him re-locate documents in any of his three displays during his ongoing investigation. In fact, while sitting at home one evening with his laptop, Noah discovered multiple links between rBear/Bert’s animal ranch and a company called “Global Ways.” Although he was able to confirm that rBear had been purchasing many endangered animals through his Global Ways connection, he could not prove that any illegitimate smuggling/reselling operations were taking place. Noah then asked his colleague, Lena, if she had any other useful information.

A short time later, Lena brought her laptop to Noah’s office, which now included several relevant documents she had found. Due to the obvious visual and annotated links that Noah had created across his three displays, Lena was able to easily catch up on the

progress of Noah’s case. The two analysts then went to work checking on the co-occurrences of entities of interest through cross-display visual links among all documents on the now four displays (including Lena’s laptop). As a consequence of the shared links, they were able to confirm that Bert and Global Ways were affiliated with a notorious illegal seller of endangered animals from Africa, and were in fact reselling them—often to private, illegal zoos or “cash-for-kill” ranches in Texas.

While this scenario represents just one example of the potential of visual links in a sensemaking documents analysis, it is emblematic of the potential of a display ecology involving collaborating users and multiple displays.

6.2 User Study

In order to evaluate the effectiveness of our SAViL tool, we conducted a qualitative human-subjects experiment. The main goal of this evaluation was to determine whether SAViL helped users create the semantic structure and synthesize their hypotheses using a broader spectrum of screens. Thus, we investigated how SAViL influenced the analysis process to enable the users to form semantic structures. This evaluation was guided by the following research questions:

- Do cross-display links help users utilize different types of displays as an integrated sensemaking space?
- Do cross-display links help users forage for and guide their attention to information on different displays?
- How is the sensemaking process different with and without cross-display linking?

This comparative study extends a number of prior related sensemaking studies describing the value of space for sensemaking, featuring large high-resolution displays [117], [23], multiple small mobile displays [24], and one created from notecards on the table [21]. In order to assess study outcomes, we investigated how the final hypotheses and distinct plots were synthesized and represented using both the visual links and multiple displays. Although we focused on analyzing and reporting observations from a cross-display link

(CL) group, we compared the impact of utilizing cross-display visual links on the analytical process and resulting product with a baseline group—the non-cross-display link group (NCL)—to determine if there were any notable differences in the sensemaking process. A summary of the comparison of the two groups is shown in Table 6.1.

Table 6.1. Evaluation results.

Group	User	Open Documents				#screens used	#screens for synthesis	#distinct plots	#plots across two displays
		Laptop	Tiled Display	TV	iMac Table top				
NCL	U1	0	11	6	0	2	1	2	1
NCL	U3	0	32	14	0	2	1	3	0
NCL	U4	0	24	0	0	1	1	2	0
NCL	U5	6	24	11	0	3	2	1	1
CL	U2	3	17	10	0	3	2	2	2
CL	U6	9	8	6	8	4	2	2	2
CL	U7	4	20	5	2	4	2	3	2
CL	U8	0	37	8	3	3	3	3	1

6.2.1 Participants

We recruited eight undergraduate participants from a local university (identified anonymously herein as U1 through U8). All eight participants were junior- and senior-level computer science majors, ranging in age from 20 to 23. Each participant verbally expressed confidence in his or her ability to solve analytical tasks; the only difference between the two groups was whether they could utilize cross-display links. Thus, the two student groups were divided as follows:

The non-cross-display link (NCL) group could not use cross-display links, but were able to use all other features, including links *within* each display.

The cross-display link (CL) group could use the cross-display link features and all other system features.

6.2.2 Dataset and Task

The experimental protocol we utilized is based on the prior work of Andrews et al. [1], Robinson [21], and Wigdor [24]—the principal difference being that we employed a display ecology. The main task for this study asked participants to perform a documents

analysis involving a collection of 63 short (up to 200 words) fictitious textual documents, requiring no special expertise or prior knowledge. Each participant performed the task individually with four displays. These text documents provided evidence of three fictitious terrorist plots and possible associated subplots, which participants were asked to identify. The participants had to overcome the critical challenge of weeding out irrelevant information on their way to identifying the fictitious plots and subplots. All participants were given several pages of letter-sized paper and pencils for note-taking.

6.2.3 Apparatus

Participants were provided with a display ecology consisting of four different display types (Figure 6.1 top):

- 60-inch tiled LCD screen (2x4 tiles with total resolution of (5120x2160) on a Windows 7 PC;
- 27-inch Apple iMac (2560x1440, laid horizontally) with a resistive touchscreen, OSX;
- 45-inch HDTV (1280x720 resolution), Windows 7 PC;
- 15-inch laptop (1366x768 resolution), Windows 7.

The participants could use any display they wanted to start the process and they were able to use the same mouse and keyboard for all four displays via an input sharing tool called Synergy [118]; alternatively, they could use a separate mouse and keyboard for each. The iMac also supported touch input. All displays were connected to independent computers. The laptops could be moved in the room but the other computers were locationally fixed. SAViL's role in this task was to provide users with enhanced tools for exploring documents in support of the analytical process.

6.2.4 Procedures

Participants first completed an informed consent form and a pre-study demographics questionnaire. Participants were then given a 10-minute tutorial on how to use the system, which focused strictly on system features. After the completion of the tutorial, participants engaged in the actual experimental session of identifying fictitious terrorist plots and subplots. After they completed the 90-minute session, the experimental session

concluded with an individual interview and survey during which the participants were encouraged to use the analytical results on the displays to support their answers.

6.2.5 Data Collection and Analysis

Throughout the session, screenshots of all of the displays were captured every 30 seconds; additionally, video recordings were utilized to capture each study session. Participants were also asked to submit any hand-written notes they took. During the post-study interviews, we asked participants from both groups to describe their findings in this order: (1) to explicitly describe the plots/subplots they identified, and (2) to identify the text documents they found to be related to any of those plots/subplots. The information from the questionnaire and the interviews allowed us to identify clearly groups of related text documents that had contributed to the formation of any plots/subplots. It should be noted, however, that our interest was principally to understand how our cross-display visual links helped users to perform the entire sensemaking process and better leverage space from multiple displays. We analyzed the collected data from a mixed-method analysis approach combining qualitative and quantitative observations.

6.3 Observations

As shown in Table 6.1, we compared outcomes based on which displays they used, how the plots were organized on the displays, and whether the hypotheses differed significantly between the CL and NCL groups.

6.3.1 Visual Link Usage Observations

An important finding is that the use of SAViL's cross-display links played a principal role in the analytical process, which was evidenced by the fact that all eight participants in the CL group were observed using some form of linking after opening a certain number of documents. This observation was true even for the NCL participants who were unable to link to other devices. Although each user employed a different analytical process and task sequence, our observations and interview results indicate that every participant used visual links to quickly identify important documents. For example, after organizing documents based on different entity types (e.g., people and places) on different displays, the

participants would often connect the visual links when the documents included keywords or entities that looked promising, in order to be able to reference those documents more easily during subsequent data analysis.

Manual linking. As corroborated by observations and interviews, five out of eight participants (U1, U2, U3, U5, and U8) used manual linking somewhat consistently. If there was information that was related to another document—but the documents did not share the same entities (e.g., an alias)—they could also be manually linked. Other participants used manual linking to tie together pieces of information that were semantically similar or represented alias-type terms (e.g., C-4 and explosive). Additionally, U3 used this feature to “shortcut” links between related documents; i.e., he would make immediate links between documents that were not otherwise automatically linked. Interestingly, some participants also linked documents that contained contradictory information and referred back to them when additional data were uncovered.

Automatic linking. The CL participants noted that this cross-display links enabled them to maintain connections between documents scattered on multiple screens, which later aided them in rapidly navigating between the related documents from different displays. For example, U2 utilized cross-display linking to see explicit connections between two documents located on different screens. U2 had previously organized several documents on different displays based on (1) geographic location and (2) people of interest. Thus, cross-display linking enabled him to further his investigation by connecting a document related to a person to another related to a place across displays.

6.3.2 Information Foraging and Awareness

The cross-display links also helped the CL participants maintain awareness of connections between documents on different displays by visually reminding them which documents were linked. Not surprisingly, the more documents they opened, the harder it was for them to locate a specific document; thus, CL users relied on visual links to locate and return to documents of interest on various displays. As U7 mentioned, “*After checking*

other documents on one screen, links make it easier to jump back to the original screen I was working on and refresh my thought process.” Three CL users (U6, U7, and U8) also mentioned that cross-visual links enabled them to separate relevant topics on different displays because the links visually illustrated how documents on different displays were tied together. This characteristic represents an important difference between the NCL and CL users (Table 6.1). As an example, U8 (a CL user) stored documents in one display based on user-linked entity type (in this case, phone numbers), and then linked them to different documents on other displays that he has previously targeted based on recurring names. Specifically, if a keyword had more links with a specific display (storing different information types), the user was able to determine quickly that the keyword had a significant relationship with a specific entity type. U8 made this observation about using visual links in foraging for information, “*Somewhat similar to reading a book, the links allow us to ‘turn page’ and keep reading from where the document left off (or just elaborate on specific details) among displays...*” This finding indicates that the ability to create cross-display visual links will assist analysts in targeting documents located on different displays, thereby significantly enhancing cross-device foraging tasks.

6.3.3 Help Leveraging Multiple Displays

We also investigated when and why a participant would add an additional screen to enhance the analytical task. When questioned during post-session interviews, two of the reasons cited by both groups were (1) to avoid the visual clutter created by too many open documents or visual links, and (2) to have more space for document organization.

Although both groups incorporated additional displays for similar reasons (space issues), other motivations also emerged. Importantly, three of four CL group members (U2, U7, and U8) tried to connect and extend their hypotheses across displays, sometimes by adding a display “hub.” For example, U2 stated the following: “*I started using another display to act as the hub connecting some hypotheses that had previously been on the tiled display.*” In addition, NCL group members U1 and U5 added new displays to start investigating separate plots: “*I wanted to differentiate between subplots on different screen*”

(U1). In short, CL participants tended to regard multiple displays supported by cross-display links as more continuous space.

Observations and post-study interviews showed that all four participants from the CL group actively used more displays in comparison to the NCL group (Table 6.1). Two of the four CL participants used all four displays to organize their documents, and the remaining two participants used three displays. As illustrated in this table (iMac table), three of the four CL users had opened at least two text documents on the iMac tabletop, while none of the NCL group had any interaction with it (Table 6.1). In post-study interviews, all four NCL participants cited poor visibility of the tabletop display. The CL participants also mentioned that the tabletop display was out of their immediate sight; nonetheless, due to the ability to utilize cross-display links they found themselves more aware of the devices around them and hence they opted to also use tabletop system. CL U8's tabletop use was directly motivated by cross-display links and the spatial position and angle of the tabletop: *"I found myself having too many links between the TV and tiled displays at the end (specifically horizontal links becoming intermingled); by moving many documents onto iMac tabletop, I could create more vertical links that were easier to distinguish."* This example represents the productive ways that CL users employed spatial organization across displays to support sensemaking tasks.

6.3.4 Semantic Structures across Displays

We compared the results of the spatial organization strategies used by both groups in order to show how the availability of cross-display visual links changed the ways they externalized information. Our observations from the spatial organization strategies employed by the NCL group resonate with observations concerning sensemaking and external memory using single large displays [1]. Overall, their spatial organization of documents was confined to the single large display space (mostly on tiled display). Specifically, NCL users appeared to have tighter clusters within a single display. As shown in Table 6.1, the number of plots the NCL group displayed across more than one display was limited to just two. The NCL participants performed linking between documents within a display (mostly tiled display), after which they clustered documents

with links closer together. Their organizational strategies using a tiled display was also influenced by the bezel size of the screens, which ultimately resulted in a grid layout. Also, some of the users performed within-display clustering to create several small organizational structures of similar documents, such as timeline and geographical locations.

Conversely, CL users preferred to use the multiple displays as if they were a single large display rather than multiple discrete entities. Even though the CL participants created similar semantic layers within each display, they were able to develop broader hypotheses by employing spatial relationships using multiple displays. The most interesting case was U6 (Figure 6.8). He used the physical location of the four screens to organize the documents in a semantic way. Specifically, he organized documents based on chronological order, effectively building a timeline of events with the earliest events on the far-left screen (laptop screen), and then progressing to the rightmost screen (the tabletop device) for more recent events. Similarly, two other CL group members (U2 and U8) used visual links to guide their organizational approach in that both divided their workspaces into different categories—suspected perpetrators on one display, places and weapons on two others. Both participants then connected these displays through the cross-display links using different topic keywords to try to gain insights into the relationships and patterns among people, places and weapons from different displays.

6.3.5 Synthesis across Displays

We also observed the way users in both groups eventually brought the information they had synthesized together after they had organized it in various ways across displays. Notably, we observed marked differences in the ways users from the two groups synthesized information from the displays. In terms of final outcomes, all of the participants were able to identify at least one or more terrorist plots. Interestingly, however, they synthesized the information pertaining to those plots differently between CL and NCL participants. For instance, three of the four NCL participants (U1, U3, and U4) reported that they were able to identify plots using only the information provided by the documents on the tiled screen.

All four participants from the CL group and U5 from the NCL group identified at least one plot as a result of utilizing documents on more than one screen. For instance, U8 (a CL user) initially organized documents across different displays based on entity types. He read the organized documents on each display and then connected interesting keywords (entities) by clicking on them in the various documents. Since the links explicitly defined the relationships, he did not have to move relevant documents from one display to another to create a tighter semantic structure.

For instance, the plot that U8 identified was determined by linking documents across displays. Specifically, U8 uncovered his “explosion plan” by linking documents on three different screens: five documents from the HDTV, two documents from the tabletop, and two documents from the tiled display (Figure 6.9).

Our observations indicate that while the usage of visual linking was consistent among all of the participants from both groups, the CL participants tended to use more screens for the process of identifying plots through multiple visual links. In contrast, three of the four participants from the NCL group used a single screen (the tiled screen) to conduct the plot identification task (Table 6.1). Although two of the NCL users attempted to synthesize information across two displays, they were hampered by their inability to employ cross-display links. We observed that these NCL participants used a significant amount switching between different displays—first checking the context of plots in documents on one display, and then referencing the same entities in other documents on another display.

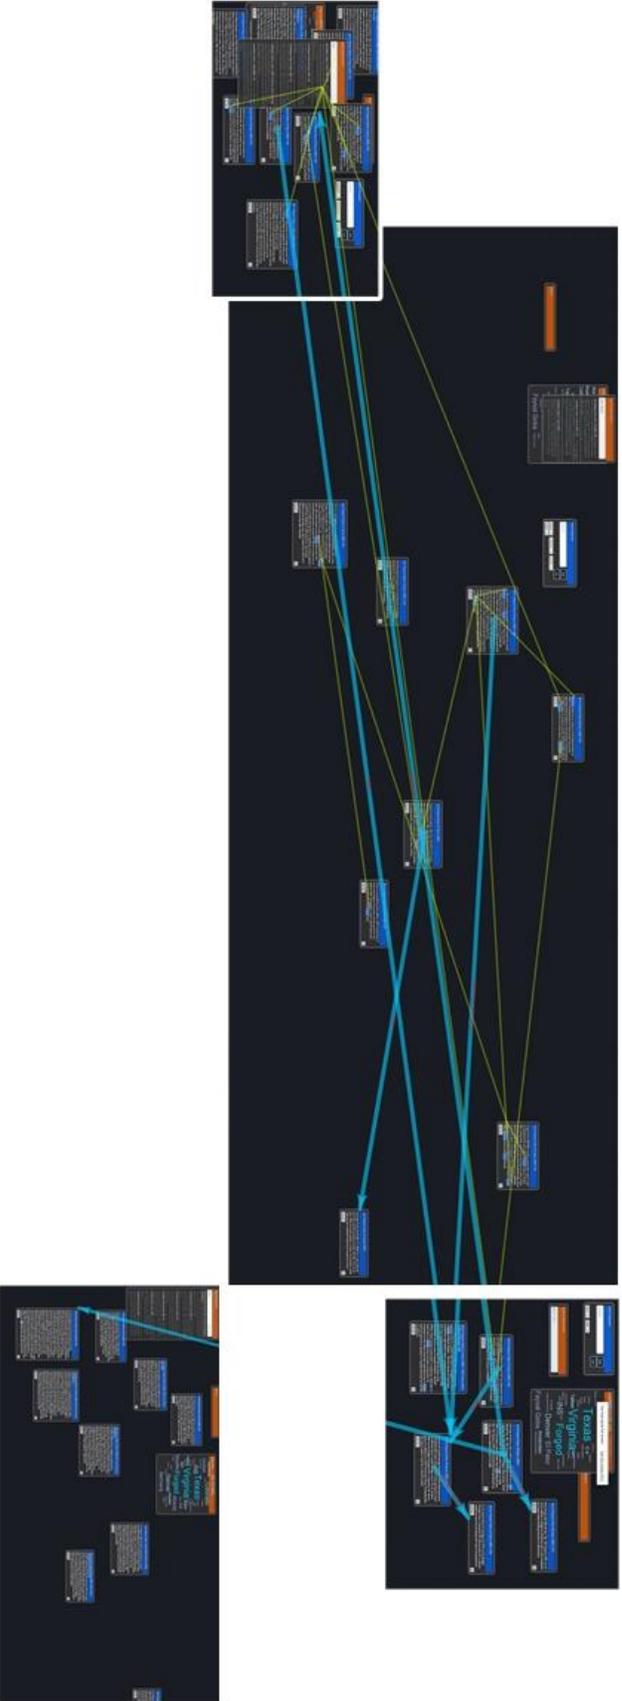


Figure 6.8: U6 organized documents based on chronological order, effectively building a timeline of events with the earliest events.

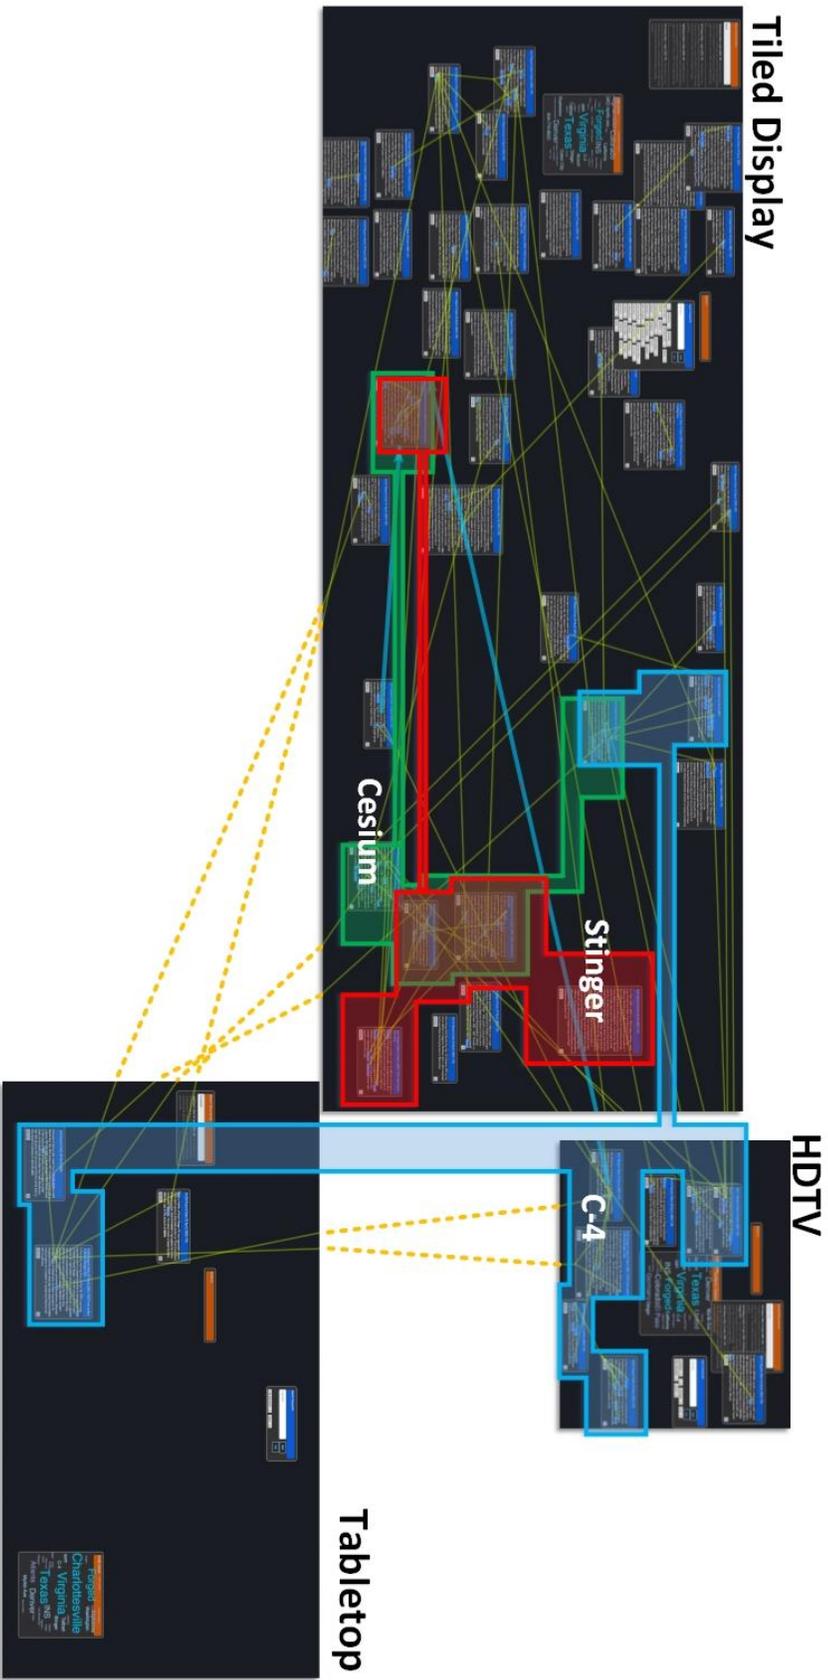


Figure 6.9. Three different plots across the displays. We added labels pertaining to participant explained different subplots and clustered documents across displays, which were described to us during the debriefing. The different colored regions represent different.

6.4 Discussion

The results of our observations and the post-session interviews show that cross-display visual links can provide new opportunities for sensemaking when analyzing text documents in a multi-display environment. Based on our results, we revisit the questions that helped guide the development of this study:

Do cross-display links help users utilize different types of displays as an integrated sensemaking space?

Throughout our evaluation, we observed that the CL users were able to extend their analysis space to different displays increasingly based on their specific needs. The cross-display visual links helped users perceive and determine space for analysis continuously. As a result, the cross-display visual links helped the CL users utilize more displays, which in turn enabled them to spatially organize and spread documents across the displays to externalize their thought processes.

Do cross-display links help users forage for and guide their attention to information on different displays?

While participants in the CL group used multiple displays to corral varying types of information, we observed that the cross-display visual links helped to direct their focus to any keyword or document located on any display at any given time. In other words, these visual links enhanced a user's cross-display foraging and re-finding tasks. Additionally, the better awareness of other displays and their information enabled users to synthesize information across different displays.

How is the data synthesis/sensemaking process different with and without cross-display linking?

In short, the CL participants formulated plots/subplots as a result of synthesizing information from documents scattered across multiple displays. This finding shows that their information synthesis was not confined within a single screen. In contrast, three of the NCL participants determined their plots/subplots using a single screen—even though additional screens were available to them. This observation shows that cross-display links

could potentially change the ways in which an analyst conducts a task using a display ecology.

Although our observations showed the positive analytical potential of utilizing cross-display visual links in performing sensemaking tasks within a display ecology, it also raised the unwanted problem of visual clutter—i.e., too many cross-display visual links. Despite the fact that we employed various link-bundling approaches, a sizable number of cross-visual links among documents and across displays still resulted in visual clutter and hindered users from viewing clear relationships. This finding confirmed the need to address this potential visual hindrance more efficiently in order to alleviate the occlusion of information. [119], [93], [14], [92]. A follow-up study is under consideration to support a novel strategy for reducing the visual clutter caused by many cross-display links.

The advantages of employing one or more visual links in a display ecology has been corroborated in that users benefited from the increased display real estate by being able to easily connect information artifacts from different displays. It should be noted, however, that our study used a simple spatial display topology, meaning that every display was proximally located—specifically within about 15cm (overlapped) to 1m. However, real-life applications will inevitably introduce the likelihood of variable display topologies. For example, a user may be working with a display topology that completely surrounds the user. Such spatial variations may become problematic if the cross-display links become overlaid in depth. In short, the spatial arrangement of a display ecology could confound a user's understanding of the visual links if they become indiscernible. Therefore, this potential problem suggests the need for future studies to investigate the implications and impact of the different display topologies.

In this study, we conducted an experiment comparing a single large screen and a heterogeneous display ecology. However, the heterogeneous displays in a display ecology can be more useful for a collaborative sensemaking scenario by multiple people in that it provides a mixture of private and public space. It appears that this work would be greatly improved by adding collaborative scenarios to the experiment.

6.5 Summary

The phrase display ecology often implies the composition of various devices in different physical locations, which has the potential to provide increased screen space, which in turn can improve performance on analytic tasks. With SAViL, our goal is to unite multiple displays into an analytic space that can be perceived as a single cohesive environment via spatially aware visual links. We conducted a qualitative study in order to evaluate the efficacy of the cross-device links feature for sensemaking tasks with display ecologies. We observed that the system helps users organize information across different displays in order to externalize and synthesize the data. Results of the study lead us to believe that our cross-display links do indeed change the way multiple displays are perceived. The participants who employed the cross-display visual links spatially organized documents, analyzed them across displays, and built hypotheses in ways that were different from the rest of the participants who did not use the feature. Participants who used the cross-device link feature tended utilize more displays and screen space to perform organization and analysis of documents. This leads us to believe that spatially-aware visual links are a critical component for transforming multiple displays into a display ecology.

7 Conclusion

This dissertation has identified important components for visual analysis in display ecologies. Our results contribute new techniques and systems for enabling visual analysis with display ecologies. This chapter summarizes the contributions made by these systems and evaluations, and then concludes with future work guided by this effort.

7.1 Restatement of Contributions

The central problem addressed by this dissertation is how to design visual analysis tools that can unite multiple displays into a single cohesive analysis space. Toward this goal, we first explored design considerations (Chapter 3) of display ecologies for visual analysis by distilling results from prior research in visual analytics, information visualization, sensemaking, distributed cognition, and multiple displays environments. Based upon our findings, we detail four essential aspects of visual analysis in a display ecology: (1) combining displays with different relationships in an analysis space, (2) transferring information among displays, (3) synthesizing information, and (4) dynamic display

memberships. These four design considerations for display ecologies represent the foundation for successful visual analysis using a display ecology. We applied these design considerations to the design of two visual analysis tools (VisPorter and SAViL) for large document datasets.

A display ecology enables analysts to distribute data and findings around all available display space—and in so doing facilitates a form of distributed cognition for data analysis. To enable users to exploit more displays and space for analyses, we developed VisPorter (Chapter 4), a collaborative visual analysis application designed to support intuitive gesture interactions for sharing and integrating information in a display ecology. Essentially, VisPorter enhances analysis tasks (e.g., information foraging and synthesis) by enabling users to distribute information across multiple personal devices (e.g., smart phones, tablets, laptops, etc.) and larger displays (e.g., HDTV, tabletop displays, wall displays, etc.). VisPorter emphasizes providing immediacy in information sharing across devices by implementing a gesture-based interface. Specifically, the user can use “flick” and “tap-hold” gestures to transfer a piece of information piece or visualization data from one screen to another in a more direct and intuitive manner.

To better understand the efficacy of VisPorter in the collaborative analysis of a large number of documents, we conducted a laboratory study for collaborative document analysis tasks utilizing multiple displays, including large displays, desktops, and small hand-held devices. Our results confirmed that VisPorter’s gesture-based transfer features allow users to extend their workspace as necessary and externalize their cognitive processes, which they accomplished by transferring individual information or concept maps from a personal tablet to nearby available large displays. These approaches also enabled users to focus attention solely on the direct physical reference of a given piece of information (e.g., a particular document, entity, or image)—rather than focusing on the data’s nominal reference, such as filename, URL, or document ID. A key benefit of this application (in addition to the immediacy of information sharing) was observed in the greater opportunity for exploiting the spatial and physical affordances of multiple displays through collaboratively creating semantic structures over multiple displays, as well as

utilizing all available display space as external memory. These activities enable users to reduce virtual navigation in synthesizing and exploring information in display ecologies.

We also investigated how a distributed model of sensemaking, spread out over multiple displays and devices, impacts the sensemaking process for the individual and for the group (Chapter 5), and whether it provides any feasible opportunities for improving the quality and efficiency of sensemaking efforts. Our study compares the use of two display models: VisCept, which is based on a model of the individual displays with shared visualization spaces; and VisPorter, which is based on the distributed model whereby different displays can be appropriated as workspaces in a unified manner. Although the general sensemaking workflow did not change across the two types of systems, we observed that the system based on the distributed model enabled a more transparent interaction for collaborations, and allowed for greater ‘objectification’ of information. Our findings have implications for how future visual analytics systems can be designed to motivate effective collaborative sensemaking with multiple displays.

Lastly, we designed and developed cross-display visual links for display ecologies (Chapter 6). For sensemaking with multiple displays, an analyst must mentally connect and synthesize pieces of relevant information in order to generate a larger coherent story. However, the challenge associated with such synthesis tasks in a display ecology is the ability to maintain awareness of and connect scattered information across separate displays, since most displays will likely be out of the user’s immediate visual field. To address this issue, I developed Spatially Aware Visual Links (SAViL), a cross-display visual link technique capable of (1) guiding the user’s attention to relevant information across displays, and (2) visually connecting related information among displays. SAViL visually represents the connections between different types of information elements (e.g., keywords, documents, pictures, etc.) across displays. Using its cross-display link feature, SAViL also enables the user to emphasize spatial relationships between displays and the physical location of displays and their information objects. To evaluate the system, I conducted a controlled user study to evaluate the impact of dynamic visual linking on sensemaking tasks for intelligence analysis in display ecologies. Participants who

employed SAViL's cross-display link feature tended to utilize more displays and screen space to perform their visual analysis tasks.

When considered collectively, the design considerations for visual analysis in a display ecology, the various visualization and interaction techniques, and the related systems described in this dissertation significantly enhance our understanding of how to accomplish visual analysis in a display ecology. In particular, one of the most notable contributions of this investigation is that the interaction and visualization techniques described herein make it possible for a display ecology to offer the same benefits of "space-to-think" [1] as large high-resolution displays—but with a significantly reduced price-tag since users will be able to combine readily accessible displays around the workspace. We have also showed that users can employ the multiple discretized screen space supported by the presented ecology systems and features as external memory (Section 4.4.3) and a variety of semantic structures (Section 4.4.4 and Section 6.3.4). We believe that this work provides users with critical components to analyze and synthesize a large amount of information via the use of a display ecology.

7.2 Limitations and Future Work

Throughout this dissertation, we have acknowledged that the system features presented herein have several limitations. In this section, we discuss these limitations, as well as potential avenues for future research suggested by those limitations.

7.2.1 The Studies

In this research, we identified several difficulties and limitations over the course of evaluating the sensemaking process with a display ecology. Indeed, one striking challenge of this research was to create effective evaluation strategies for human sensemaking in the context of a display ecology. First of all, a real-life sensemaking scenario is highly unpredictable and may not have one specific solution. Analysts selectively encounter, consult, and retain various pieces of information at opportunistic moments, transitioning between spaces throughout the day, the week (or longer) as needed. Moreover, the amount of information and data pertaining to any given scenario is virtually limitless.

Therefore, a longitudinal study may be more appropriate for better understanding the complex characteristics of sensemaking using a display ecology. In contrast, the participants who took part in our studies were given a clear goal within a controlled lab analysis setting. Our dataset was comprised of a relatively small number of documents—which is vastly different from the essentially inexhaustible amount of data used in authentic intelligence analysis scenarios. While these issues may have reduced the ecological validity of the study somewhat, the inevitable restrictions we faced (i.e., time and financial considerations) required a more feasible analysis task that (1) could be completed within 90 minutes, and (2) used a manageable dataset. I speculate that when tested in a longitudinal setting, the benefits of a display ecology will become even more apparent.

Also, the evaluation and user studies of the ecology systems featured in this investigation focused primarily on (1) showing qualitatively how users externalize their thought/sensemaking processes with multiple displays, and (2) how our presented systems and techniques impact the strategy and process of visual analysis. Specifically, we evaluated how multiple displays enable the creation of a more powerful “space to think,” whereby users can employ the discretized screen space to spatially organize information and data elements across different displays. Based on user scores and feedback, we determined that the “space-to-think” activities represent the most important factors in performing sensemaking tasks in display ecologies—principally because they enable users to better exploit human spatial senses and the physical space facilitated by display ecologies for enhanced analysis [21], [1]. However, in addition to such analysis activities, it was clear that one must acknowledge the variability and diversity of analysis styles among different users, which inevitably affect the sensemaking process and performance.

Another limitation that must be noted is that our studies were based on fixed types of displays; in contrast, heterogeneous displays dynamically chosen by users would likely affect their methods and strategies for information gathering, thereby impacting the overall sensemaking performance. Thus, in a future user study I will investigate and

illustrate the pros and cons of various analysis patterns and dynamic display ecologies with the help of our presented systems and techniques.

In addition, even though we discussed the following issues in Section 4.5.2, we need to further clarify and evaluate the decision-making processes involved in a user's preferred analysis strategy. Specifically, how and why do users decide on their initial analysis strategies and appropriate a particular set of heterogeneous displays? What factors influence those decisions? And finally, what user analysis strategies appear to be the most effective? Understanding more diverse display ecologies and analysis strategies will offer new insights and implications for designing novel visual analysis tools.

In our studies of sensemaking and multi-display usage, we determined that it can be difficult to appropriately attribute actions to motives. For example, in the VisPorter study, the document flicking actions can potentially embed a variety of meanings based on a user's intentions—for example, offloading, self or team referencing, or simple transfer between displays. Similarly, it is still relatively unclear what motivates users to move documents between and among displays. To understand their motivations and identify preferred analysis patterns, we depended fully on post-study interviews and quotes—but this approach has shown limitations since their statements may not have reflected their true motivations. A future study could utilize a “think aloud protocol” to ask the participants about their intentions when they are flicking documents or conducting other tasks within a display ecology.

Lastly, the two studies of my display ecology systems (detailed in Chapters 4 and 6) have different limitations related to the social relationships among study participants, which could have impacted their findings. In the VisPorter study, for example, the social relationships among the study's cohort were minimally considered in evaluating their sensemaking decisions and strategies. In other words, whether or not the participants knew and/or trusted each other (and to what degree) may have significantly affected their collaboration styles and performance. Thus, a future study should compare the

sensemaking strategies employed by participants who know and trust one another with users who do not.

7.2.2 Automatic View Adaptation for Multiple Displays

The visual representation and content of relevant information should be adapted based on the properties of the different displays and the user's preferences and needs. Throughout the display ecology study sessions, many participants asked if they could alter the size of documents and fonts depending on the display size. Re-rendering based on the display properties is required when visual information or visualizations are transferred from one display to different types of displays. When a visualization is exhibited on various displays, the visual representation and content of the visualization must be resized and adapted according to the properties of the displays automatically.

For analysis and presentation of data visualizations in general, resizing is particularly critical in the context of a dashboard with limited real-estate, and/or when visualizations created on one display must then be rendered on a different display. The major challenge associated with techniques supporting resizing and creating multi-scale visualization is the significant number of variations that must be considered. In fact, it is almost impossible for a visualization designer to consider every possible combination of display resolution, size, and aspect ratio. SAViL (Chapter 6) partially addressed this challenge by automatically adjusting the link thinness and the font size, based on the pixel density of different displays—but it is crucial for visualization techniques to support a smarter way to automatically adapt and represent more complex visualization view based upon different scales.

Inspired by the principle of cartographic generalization, I will explore smarter ways to adapt and simplify a visualization based upon different scales (Figure 7.1). I will investigate optimization techniques for resizing visualizations based on the spatial constraints and semantics of the visualization view, both of which inform the level of detail rendered at a given scale and on different types of displays.

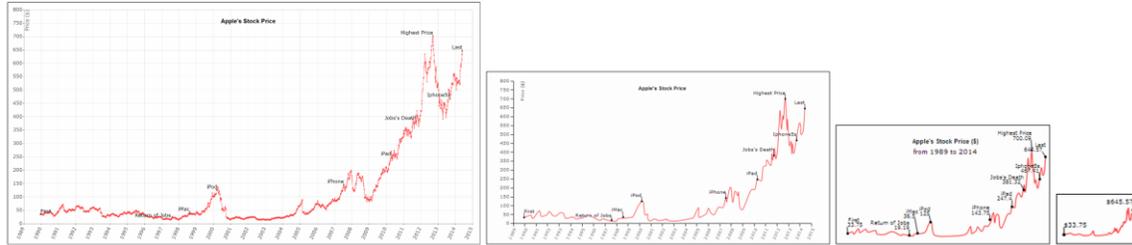


Figure 7.1. An example of a line chart automatically rendered at different scales. The algorithm preserves the various elements of the line chart based on their semantic importance at a given scale.

7.2.3 Support for Software Framework and Infrastructure

Building visual analysis tools based on multiple heterogeneous devices is very difficult due to system-imposed constraints, such as the heterogeneity of communication protocols and different software and hardware platforms. However, separate displays should be able to easily communicate with each other for analysis tasks, thus allowing users to employ any nearby display as an extension of the devices they are currently using when and as needed in the analysis context. In short, it is essential that a display ecology infrastructure support this interoperability. For visual analysis in a display ecology, flexible interoperability requires several important capabilities, including: (1) information transfer between devices, (2) spatial co-awareness between devices, (3) linking multiple device displays into a common underlying information space, (4) the use of one device as an interaction input for another display device, and (5) dynamic device membership in ecologies.

In response to these essential challenges and requirements, I will construct a software framework that is based on two primary contributions: (1) a web-based infrastructure in which the information and user events from different displays can be distributed and synchronized across different computing devices, and (2) an easy-to-use programming toolkit that supports a set of reusable interactions and visualization techniques spanning multiple displays and devices.

7.2.4 Analysis Provenance for Display Ecologies

Finally, the systems presented in this dissertation lack support for provenance [120], which might hamper an analyst’s full use of the space. In our studies, we noted that many participants were concerned that they might lose information when multiple collaborators were moving information among multiple displays. Our presented systems provided a very high degree of freedom in spatially organizing and distributing information across different devices and displays. Thus, it was challenging for users to keep track of changes made. The provenance of information can help users understand how their analytical steps using multiple devices derived a final hypothesis—such as IdeaVis’s Facilitator display, which provides information relating to the work process and history for collaborative sketching sessions [90].

Final Remarks

Our current computing environment requires new ways for leveraging a large number of available displays to explore and analyze large, complex data aggregates. In this dissertation, we argue that a display ecology enables users to exploit the burgeoning interaction opportunities for visual analysis, which are made possible by the modern technological landscape—one where most people possess multiple computational and interactional resources such as laptops, smartphones, and tablets. We hope that this dissertation guides the design of new visualizations and visual analytics systems for display ecologies and presents inspiration for future research in ubiquitous analysis scenarios.

References

- [1] C. Andrews, A. Endert, and C. North, "Space to think: large high-resolution displays for sensemaking," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 55-64, 2010.
- [2] S. Klum, P. Isenberg, R. Langner, J.-D. Fekete, and R. Dachzelt, "Stackables: combining tangibles for faceted browsing," in *Proceedings of the International Working Conference on Advanced Visual Interfaces*, pp. 241-248, 2012.
- [3] J. C. Lee, S. E. Hudson, J. W. Summet, and P. H. Dietz, "Moveable interactive projected displays using projector based tracking," in *Proceedings of the ACM symposium on User interface software and technology*, pp. 63-72, 2005.
- [4] S. K. Card, J. D. Mackinlay, and B. Shneiderman, *Readings in information visualization: using vision to think*. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 1999.
- [5] E. M. Huang, E. D. Mynatt, and J. P. Trimble, "Displays in the wild: understanding the dynamics and evolution of a display ecology," in *Proceedings of the 4th international conference on Pervasive Computing*, pp. 321-336, 2006.
- [6] P. Isenberg, T. Isenberg, T. Hesselmann, B. Lee, U. Von Zadow, and A. Tang, "Data Visualization on Interactive Surfaces: A Research Agenda," *IEEE Computer Graphics and Applications*, vol. 33, 2013.
- [7] J. Hollan, E. Hutchins, and D. Kirsh, "Distributed cognition: toward a new foundation for human-computer interaction research," *ACM Transactions on Computer-Human Interaction*, vol. 7, pp. 174-196, 2000.
- [8] H. Chung, Y. Seungwon, N. Massjouni, C. Andrews, R. Kanna, and C. North, "VizCept: Supporting synchronous collaboration for constructing visualizations in intelligence analysis," in *Proceedings of IEEE Visual Analytics Science and Technology (VAST) 2010*, pp. 107-114, 2010.
- [9] K. Krippendorff, *The semantic turn: A new foundation for design*: crc Press, 2005.
- [10] A. Crabtree and T. Rodden, "Hybrid ecologies: understanding cooperative interaction in emerging physical-digital environments," *Personal and Ubiquitous Computing*, vol. 12, pp. 481-493, 2008.
- [11] B. A. Nardi and V. O'Day, *Information Ecologies: Using Technology with Heart*: MIT Press, 1999.
- [12] T. Coughlan, T. D. Collins, A. Adams, Y. Rogers, P. A. Haya, and E. Martín, "The conceptual framing, design and evaluation of device ecologies for collaborative

- activities," *International Journal of Human-Computer Studies*, vol. 70, pp. 765-779, no. 10, 2012.
- [13] L. Terrenghi, A. Quigley, and A. Dix, "A taxonomy for and analysis of multi-person-display ecosystems," *Personal Ubiquitous Comput.*, vol. 13, pp. 583-598, 2009.
- [14] M. Waldner, A. Lex, M. Streit, and D. Schmalstieg, "Design considerations for collaborative information workspaces in multi-display environments," in *Proceedings of Collaborative Visualization on Interactive Surfaces-CoVIS'09*, 2009.
- [15] S. K. Badam and N. Elmqvist, "PolyChrome: A Cross-Device Framework for Collaborative Web Visualization," in *Proceedings of the ACM International Conference on Interactive Tabletops and Surfaces*, pp. 109-118, 2014.
- [16] W. Javed and N. Elmqvist, "Exploring the design space of composite visualization," in *Proceedings of IEEE Pacific Visualization Symposium (PacificVis) 2012*, pp. 1-8, 2012.
- [17] R. Ball, C. North, and D. A. Bowman, "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*, pp. 191-200, 2007.
- [18] K. Vogt, L. Bradel, C. Andrews, C. North, A. Endert, and D. Hutchings, "Co-located collaborative sensemaking on a large high-resolution display with multiple input devices," in *Proceedings of Human-Computer Interaction-INTERACT 2011*, pp. 589-604, 2011.
- [19] M. Levin, *Designing Multi-device Experiences: An Ecosystem Approach to User Experiences Across Devices*: " O'Reilly Media, Inc.", 2014.
- [20] C. North, "Visualization Viewpoints: Toward Measuring Visualization Insight," *IEEE Computer Graphics & Applications*, vol. 26, pp. 6-9, 2006.
- [21] A. C. Robinson, "Collaborative synthesis of visual analytic results," in *Proceedings of IEEE Visual Analytics Science and Technology 2008*, pp. 67-74, 2008.
- [22] C. Andrews and C. North, "The Impact of Physical Navigation on Spatial Organization for Sensemaking," *Visualization and Computer Graphics, IEEE Transactions on*, vol. 19, pp. 2207-2216, 2013.
- [23] C. Andrews and C. North, "Analyst's Workspace: An embodied sensemaking environment for large, high-resolution displays," in *Proceedings of IEEE Visual Analytics Science and Technology 2012*, pp. 123-131, 2012.
- [24] P. Hamilton and D. J. Wigdor, "Conductor: enabling and understanding cross-device interaction," in *Proceedings of the SIGCHI conference on Human factors in computing systems*, pp. 2773-2782, 2014.

- [25] W. Wright, D. Schroh, P. Proulx, A. Skaburskis, and B. Cort, "The Sandbox for analysis: concepts and methods," in *Proceedings of the SIGCHI conference on Human factors in computing systems*, pp. 801-810, 2006.
- [26] J. Stasko, C. Gorg, and Z. Liu, "Jigsaw: supporting investigative analysis through interactive visualization," *Information visualization*, vol. 7, pp. 118-132, 2008.
- [27] i2Limited. (2010). *Analyst's Notebook*. Available: <http://www.i2inc.com>.
- [28] E. A. Bier, E. W. Ishak, and E. H. Chi, "Entity workspace: an evidence file that aids memory, inference, and reading. ," in *Proceedings of IEEE International Conference on Intelligence and Security Informatics (ISI 2006)*, 2006.
- [29] E. A. Bier, S. K. Card, and J. W. Bodnar, "Entity-based collaboration tools for intelligence analysis," in *Proceedings of IEEE Visual Analytics Science and Technology (VAST)*, pp. 99-106, 2008.
- [30] W. A. Pike, J. Bruce, B. Baddeley, D. Best, L. Franklin, R. May, *et al.*, "The Scalable Reasoning System: Lightweight visualization for distributed analytics," *Information Visualization*, vol. 8, pp. 71-84, no. 1, 2008.
- [31] F. B. Viegas, M. Wattenberg, F. v. Ham, J. Kriss, and M. McKeon, "ManyEyes: a Site for Visualization at Internet Scale," *IEEE Transactions on Visualization and Computer Graphics*, vol. 13, pp. 1121-1128, 2007.
- [32] M. McKeon, "Harnessing the Information Ecosystem with Wiki-based Visualization Dashboards," *IEEE Transactions on Visualization and Computer Graphics*, vol. 15, pp. 1081-1088, 2009.
- [33] J. Heer, F. B. Viégas, and M. Wattenberg, "Voyagers and voyeurs: supporting asynchronous collaborative information visualization," in *Proceedings of the SIGCHI conference on Human factors in computing systems*, pp. 1029-1038, 2007.
- [34] Tableau. (2010). *Tableau Public*. Available: <https://public.tableau.com/>
- [35] P. Isenberg, A. Tang, and S. Carpendale, "An exploratory study of visual information analysis," in *Proceedings of the SIGCHI conference on Human factors in computing systems*, pp. 1217-1226, 2008.
- [36] M. Tobiasz, P. Isenberg, and S. Carpendale, "Lark: Coordinating Co-located Collaboration with Information Visualization," *IEEE Transactions on Visualization and Computer Graphics*, vol. 15, pp. 1065-1072, 2009.
- [37] K. Kim, W. Javed, C. Williams, N. Elmqvist, and P. Irani, "Hugin: A framework for awareness and coordination in mixed-presence collaborative information visualization," in *ACM International Conference on Interactive Tabletops and Surfaces*, pp. 231-240, 2010.
- [38] W. McGrath, B. Bowman, D. McCallum, J. D. Hincapié-Ramos, N. Elmqvist, and P. Irani, "Branch-explore-merge: facilitating real-time revision control in

- collaborative visual exploration," in *Proceedings of the ACM international conference on Interactive tabletops and surfaces 2012*, pp. 235-244, 2012.
- [39] H.-C. Jetter, J. Gerken, M. Zöllner, H. Reiterer, and N. Milic-Frayling, "Materializing the query with facet-streams: a hybrid surface for collaborative search on tabletops," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 3013-3022, 2011.
- [40] P. Isenberg, S. Carpendale, A. Bezerianos, N. Henry, and J.-D. Fekete, "CoCoNutTrix: collaborative retrofitting for information visualization," *IEEE Comput. Graph. Appl.*, vol. 29, pp. 44-57, 2009.
- [41] P. Isenberg, D. Fisher, M. R. Morris, K. Inkpen, and M. Czerwinski, "An exploratory study of co-located collaborative visual analytics around a tabletop display," in *Proceedings of IEEE Visual Analytics Science and Technology (VAST) 2010*, pp. 179-186, 2010.
- [42] D. Wigdor, C. Shen, C. Forlines, and R. Balakrishnan, "Table-centric interactive spaces for real-time collaboration," in *Proceedings of the International Working Conference on Advanced Visual Interfaces*, pp. 103-107, 2006.
- [43] B. Johanson, S. Ponnekanti, C. Sengupta, and A. Fox, "Multibrowsing: Moving Web Content across Multiple Displays," in *Proceedings of the international conference on Ubiquitous Computing*, pp. 346-353, 2001.
- [44] A. Fox, B. Johanson, P. Hanrahan, and T. Winograd, "Integrating information appliances into an interactive workspace," *IEEE Computer Graphics and Applications*, vol. 20, pp. 54-65, 2000.
- [45] D. S. Tan, B. Meyers, and M. Czerwinski, "WinCuts: manipulating arbitrary window regions for more effective use of screen space," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 1525-1528, 2004.
- [46] D. Wigdor, H. Jiang, C. Forlines, M. Borkin, and C. Shen, "WeSpace: the design development and deployment of a walk-up and share multi-surface visual collaboration system," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 1237-1246, 2009.
- [47] H. Jiang, D. Wigdor, C. Forlines, M. Borkin, J. Kauffmann, and C. Shen, "LivOlay: interactive ad-hoc registration and overlapping of applications for collaborative visual exploration," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 1357-1360, 2008.
- [48] S. Izadi, H. Brignull, T. Rodden, Y. Rogers, and M. Underwood, "Dynamo: a public interactive surface supporting the cooperative sharing and exchange of media," in *Proceedings of the ACM symposium on User interface software and technology*, pp. 159-168, 2003.

- [49] S. Greenberg, M. Boyle, and J. Laberge, "PDAs and shared public displays: Making personal information public, and public information personal," *Personal and Ubiquitous Computing*, vol. 3, pp. 54-64, 1999.
- [50] O. Chapuis, A. Bezerianos, and S. Frantzeskakis, "Smarties: An input system for wall display development," in *Proceedings of the ACM conference on Human factors in computing systems*, pp. 2763-2772, 2014.
- [51] Y. Jansen, P. Dragicevic, and J.-D. Fekete, "Tangible Remote Controllers for Wall-Size Displays," in *Proceedings of the ACM conference on Human factors in computing systems*, pp. 2865-2874, 2012.
- [52] M. Spindler, C. Tominski, H. Schumann, and R. Dachsel, "Tangible views for information visualization," in *Proceedings of the ACM conference on Human factors in computing systems*, pp. 157-166, 2010.
- [53] S. Volda, M. Tobiasz, J. Stromer, P. Isenberg, and S. Carpendale, "Getting practical with interactive tabletop displays: designing for dense data, fat fingers, diverse interactions, and face-to-face collaboration," in *Proceedings of the ACM International Conference on Interactive Tabletops and Surfaces*, pp. 109-116, 2009.
- [54] R. Rädle, H.-C. Jetter, N. Marquardt, H. Reiterer, and Y. Rogers, "HuddleLamp: Spatially-Aware Mobile Displays for Ad-hoc Around-the-Table Collaboration," in *Proceedings of the ACM International Conference on Interactive Tabletops and Surfaces*, pp. 45-54, 2014.
- [55] J. Cole, K. Daly, J. Haria, and P. Jain. (2012). *ScreenSquared: Improving Media Sharing with Smartphones*. Available: <http://www.prabhavjain.com/files/receiver.pdf>.
- [56] R. Borovoy and B. Knep, "Junkyard jumbotron," *MIT Center for Civic Media. Retrieved April*, vol. 15, p. 2013, 2011.
- [57] D. Schmidt, F. Chehimi, E. Rukzio, and H. Gellersen, "PhoneTouch: a technique for direct phone interaction on surfaces," in *Proceedings of the ACM symposium on User interface software and technology*, pp. 13-16, 2010.
- [58] J. Seifert, A. Simeone, D. Schmidt, P. Holleis, C. Reinartz, M. Wagner, *et al.*, "MobiSurf: improving co-located collaboration through integrating mobile devices and interactive surfaces," in *Proceedings of the ACM international conference on Interactive tabletops and surfaces*, pp. 51-60, 2012.
- [59] M. A. Nacenta, C. Gutwin, D. Aliakseyeu, and S. Subra-manian, "There and back again: cross-display object movement in multi-display environments," *Human Computer Interaction*, vol. 24, pp. 170-229, 2009.
- [60] J. Rekimoto, "Pick-and-drop: a direct manipulation technique for multiple computer environments," in *Proceedings of the ACM symposium on User interface software and technology*, pp. 31-39, 1997.

- [61] R. Dachsel and R. Buchholz, "Throw and Tilt-Seamless Interaction across Devices Using Mobile Phone Gestures," *GI Jahrestagung (1)*, pp. 272-278, 2008.
- [62] N. Marquardt, K. Hinckley, and S. Greenberg, "Cross-device interaction via micro-mobility and f-formations," in *Proceedings of the ACM symposium on User interface software and technology*, pp. 13-22, 2012.
- [63] M. A. Nacenta, S. Sakurai, T. Yamaguchi, Y. Miki, Y. Itoh, Y. Kitamura, *et al.*, "E-conic: a perspective-aware interface for multi-display environments," in *Proceedings of the ACM symposium on User interface software and technology*, pp. 279-288, 2007.
- [64] C. Pirchheim, M. Waldner, and D. Schmalstieg, "Deskotheque: Improved spatial awareness in multi-display environments," in *IEEE Virtual Reality Conference (VR 2009)*, pp. 123-126, 2009.
- [65] R. Xiao, M. A. Nacenta, R. L. Mandryk, A. Cockburn, and C. Gutwin, "Ubiquitous cursor: a comparison of direct and indirect pointing feedback in multi-display environments," in *Proceedings of Graphics Interface 2011*, pp. 135-142, 2011.
- [66] P. Baudisch and R. Rosenholtz, "Halo: a technique for visualizing off-screen objects," in *Proceedings of the SIGCHI conference on Human factors in computing systems*, pp. 481-488, 2003.
- [67] K. Hinckley, G. Ramos, F. Guimbretiere, P. Baudisch, and M. Smith, "Stitching: pen gestures that span multiple displays," in *Proceedings of the working conference on Advanced visual interfaces*, pp. 23-31, 2004.
- [68] M. A. Nacenta, S. Sallam, B. Champoux, S. Subramanian, and C. Gutwin, "Perspective cursor: perspective-based interaction for multi-display environments," in *Proceedings of the SIGCHI conference on Human Factors in computing systems*, pp. 289-298, 2006.
- [69] N. A. Streitz, J. Geißler, T. Holmer, *et al.*, "i-LAND: an interactive landscape for creativity and innovation," in *Proceedings of the SIGCHI conference on Human Factors in computing systems*, pp. 120-127, 1999.
- [70] H.-C. Jetter, M. Zöllner, J. Gerken, and H. Reiterer, "Design and Implementation of Post-WIMP Distributed User Interfaces with ZOIL," *International Journal of Human-Computer Interaction*, vol. 28, pp. 737-747, 2012.
- [71] F. Geyer, H. C. Jetter, U. Pfeil, and H. Reiterer, "Collaborative sketching with distributed displays and multimodal interfaces," in *ACM International Conference on Interactive Tabletops and Surfaces*, pp. 259-260, 2010.
- [72] C. Forlines and R. Lilien, "Adapting a single-user, single-display molecular visualization application for use in a multi-user, multi-display environment," *Proceedings of the working conference on Advanced visual interfaces*, pp. 367-371, 2008.

- [73] M. Modahl, I. Bagrak, M. Wolenetz, P. Hutto, and U. Ramachandran, "Mediabroker: An architecture for pervasive computing," in *Proceedings of IEEE Pervasive Computing and Communications 2004*, pp. 253-262, 2004.
- [74] N. Elmqvist, "Distributed user interfaces: State of the art," in *Distributed User Interfaces*, ed: Springer, pp. 1-12, 2011.
- [75] L. Frosini and F. Paternò, "User interface distribution in multi-device and multi-user environments with dynamically migrating engines," in *Proceedings of the ACM symposium on Engineering interactive computing systems*, pp. 55-64, 2014.
- [76] J. Yang and D. Wigdor, "Panelrama: enabling easy specification of cross-device web applications," in *Proceedings of the ACM conference on Human factors in computing systems*, pp. 2783-2792, 2014.
- [77] M. Nebeling, T. Mints, M. Husmann, and M. Norrie, "Interactive development of cross-device user interfaces," in *Proceedings of the 32nd annual ACM conference on Human factors in computing systems*, pp. 2793-2802, 2014.
- [78] S. K. Badam, E. Fisher, and N. Elmqvist, "Munin: A Peer-to-Peer Middleware for Ubiquitous Analytics and Visualization Spaces," *Visualization and Computer Graphics, IEEE Transactions on*, vol. 21, pp. 215-228, 2015.
- [79] P. Pirolli and S. Card, "Sensemaking Processes of Intelligence Analysts and Possible Leverage Points as Identified Through Cognitive Task Analysis," in *Proceedings of International Conference on Intelligence Analysis*, p. 6, 2005.
- [80] M. Li and L. Kobbelt, "Dynamic tiling display: building an interactive display surface using multiple mobile devices," in *Proceedings of the 11th International Conference on Mobile and Ubiquitous Multimedia*, p. 24, 2012.
- [81] A. Lucero, J. Holopainen, and T. Jokela, "Pass-them-around: collaborative use of mobile phones for photo sharing," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 1787-1796, 2011.
- [82] E. R. Fisher, S. K. Badam, and N. Elmqvist, "Designing peer-to-peer distributed user interfaces: Case studies on building distributed applications," *International Journal of Human-Computer Studies*, vol. 72, pp. 100-110, 2014.
- [83] F. Geyer and H. Reiterer, "A cross-device spatial workspace supporting artifact-mediated collaboration in interaction design," in *Proceedings of CHI '10 Extended Abstracts*, pp. 3787-3792, 2010.
- [84] F. Geyer, U. Pfeil, J. Budzinski, A. Höchtl, and H. Reiterer, "Affinitytable—a hybrid surface for supporting affinity diagramming" in *Human-computer interaction*, ed: Springer, pp. 477-484, 2011.
- [85] J. Sanneblad and L. E. Holmquist, "Ubiquitous graphics: combining hand-held and wall-size displays to interact with large images," in *Proceedings of the working conference on Advanced visual interfaces*, pp. 373-377, 2006.

- [86] B. A. Myers, R. C. Miller, B. Bostwick, and C. Evankovich, "Extending the windows desktop interface with connected handheld computers," in *4th USENIX Windows Systems Symposium*, pp. 79-88, 2000.
- [87] F. Chehimi and E. Rukzio, "Throw your photos: an intuitive approach for sharing between mobile phones and interactive tables," in *Proceedings of the 12th ACM international conference on Ubiquitous computing-Adjunct*, pp. 443-444, 2010.
- [88] D. Schmidt, J. Seifert, E. Rukzio, and H. Gellersen, "A cross-device interaction style for mobiles and surfaces," in *Proceedings of the Designing Interactive Systems Conference*, pp. 318-327, 2012.
- [89] D. Dearman and J. S. Pierce, "It's on my other computer!: computing with multiple devices," in *Proceedings of the SIGCHI Conference on Human factors in Computing Systems*, pp. 767-776, 2008.
- [90] F. Geyer, J. Budzinski, and H. Reiterer, "IdeaVis: a hybrid workspace and interactive visualization for paper-based collaborative sketching sessions," in *Proceedings of the 7th Nordic Conference on Human-Computer Interaction: Making Sense Through Design*, pp. 331-340, 2012.
- [91] Microsoft. (2009). *DeskPiles*. Available: <http://research.microsoft.com/en-us/projects/desktopiles/>.
- [92] M. Waldner, W. Puff, A. Lex, M. Streit, and D. Schmalstieg, "Visual links across applications," in *Proceedings of Graphics Interface 2010*, pp. 129-136, 2010.
- [93] M. Steinberger, M. Waldner, M. Streit, A. Lex, and D. Schmalstieg, "Context-preserving visual links," *IEEE Transactions on Visualization and Computer Graphics*, vol. 17, pp. 2249-2258, 2011.
- [94] S. Gustafson, P. Baudisch, C. Gutwin, and P. Irani, "Wedge: clutter-free visualization of off-screen locations," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 787-796, 2008.
- [95] H. Chung, Y. J. Cho, J. Self, and C. North, "Pixel-oriented Treemap for multiple displays," in *IEEE Visual Analytics Science and Technology (VAST)*, pp. 289-290, 2012.
- [96] J. Schwarz, D. Klionsky, C. Harrison, P. Dietz, and A. Wilson, "Phone as a pixel: enabling ad-hoc, large-scale displays using mobile devices," *Proceedings of the SIGCHI Conference on Human factors in Computing Systems*, pp. 2235-2238, 2012.
- [97] W. Kuhn, "Handling data spatially: Spatializing user interfaces," in *Advances in GIS Research - 7th International Symposium on Spatial Data Handling (SDH'96)*, pp. 877-893, 1996.
- [98] D. Norman, "Affordance, Conventions and Design," *Interactions*, vol. May+June, pp. 38-42, 1999.

- [99] N. Marquardt, "Proxemic interactions in ubiquitous computing ecologies," in *Proceedings of CHI'11 Extended Abstracts*, pp. 1033-1036, 2011.
- [100] F. Geyer, U. Pfeil, A. Höchtl, J. Budzinski, and H. Reiterer, "Designing reality-based interfaces for creative group work," in *Proceedings of the ACM conference on Creativity and cognition*, pp. 165-174, 2011.
- [101] K. Everitt, C. Shen, K. Ryall, and C. Forlines, "MultiSpace: Enabling electronic document micro-mobility in table-centric, multi-device environments," in *Horizontal Interactive Human-Computer Systems, 2006. TableTop 2006. First IEEE International Workshop on*, 2006, p. 8 pp.
- [102] S. L. C. Y. Yee, F. Quek, A. Endert, C. Haeyong, and B. Sawyer, "The Physicality of Technological Devices in Education: Building a Digital Experience for Learning," in *Proceedings of IEEE International Conference on Advanced Learning Technologies (ICALT)*, pp. 579-581, 2012.
- [103] C. G. Grinstein, S. Plaisant, T. Laskowski, J. S. O'Connell, and a. M. Whiting., "VAST 2007 Contest-Blue Iguanodon," in *IEEE Symposium on Visual Analytics Science and Technology*, pp. 231-232, 2007.
- [104] G. a. M. J. M. Salton, *Introduction to modern information retrieval*: McGraw-Hill, 1983.
- [105] B. Baldwin and B. Carpenter. (2003). *Lingpipe*. Available: <http://alias-i.com/lingpipe/>.
- [106] A. J. Cañas, J. D. Novak, F. M. González, A. J. Cañas, G. Hill, R. Carff, *et al.*, "CmapTools: A Knowledge Modeling and Sharing Environment," in *The first Int. Conference on Concept Mapping*, pp. 125-133, 2004.
- [107] I. Fette and A. Melnikov. (2011). *The WebSocket protocol*. Available: <http://tools.ietf.org/html/rfc6455>.
- [108] D. Crockford. (2006). *The application/json media type for javascript object notation (json)*. Available: <http://tools.ietf.org/html/rfc4627>.
- [109] S. Tilkov and S. Vinoski, "Node. js: Using JavaScript to build high-performance network programs," *IEEE Internet Computing*, vol. 14, pp. 80-83, 2010.
- [110] M. Kaltenbrunner, T. Bovermann, R. Bencina, and E. Costanza, "TUIO: A protocol for table-top tangible user interfaces," in *Proceedings of the The 6th Int'l Workshop on Gesture in Human-Computer Interaction and Simulation*, 2005.
- [111] F. Hughes and D. Schum, *Discovery-proof-choice, the art and science of the process of intelligence analysis – preparing for the future of intelligence analysis*: Joint Military Intelligence College, 2003.
- [112] C. Plaisant, G. Grinstein, J. Scholtz, M. Whiting, and e. al., "Evaluating Visual Analytics at the 2007 VAST Symposium Contest," *IEEE Comput. Graph. Appl.*, vol. 28, pp. 12-21, 2008.

- [113] F. M. Shipman III and C. C. Marshall, "Formality considered harmful: Experiences, emerging themes, and directions on the use of formal representations in interactive systems," *Computer Supported Cooperative Work (CSCW)*, vol. 8, pp. 333-352, 1999.
- [114] H. H. Clark and S. E. Brennan, "Grounding In Communication," *Perspectives on Socially Shared Cognition*, vol. 13, pp. 127-149, 1991.
- [115] A. B. Sekuler and R. F. Murray, "9 Amodal completion: A case study in grouping," *Advances in Psychology*, vol. 130, pp. 265-293, 2001.
- [116] D. Holten and J. J. Van Wijk, "Force-Directed Edge Bundling for Graph Visualization," in *Computer Graphics Forum*, pp. 983-990, 2009.
- [117] A. Endert, P. Fiaux, and C. North, "Semantic interaction for visual text analytics," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 473-482, 2012.
- [118] N. Bolton. (2014). *Synergy*. Available: <http://synergy-foss.org/>.
- [119] A. O. Artero, M. C. F. de Oliveira, and H. Levkowitz, "Uncovering clusters in crowded parallel coordinates visualizations," in *Proceedings of INFOVIS 2004*, pp. 81-88, 2004.
- [120] C. T. Silva, J. Freire, and S. P. Callahan, "Provenance for Visualizations: Reproducibility and Beyond," *Computing in Science & Engineering*, vol. 9, pp. 82-89, 2007.