

Four Considerations for Supporting Visual Analysis in Display Ecologies

Haeyong Chung*, Chris North†, Sarang Joshi†, Jian Chen‡

University of Alabama Huntsville*

Virginia Tech†

University of Maryland Baltimore County‡

ABSTRACT

The current proliferation of large displays and mobile devices presents a number of exciting opportunities for visual analytics and information visualization. The display ecology enables multiple displays to function in concert within a broader technological environment to accomplish visual analysis tasks. Based on a comprehensive survey of multi-display systems from a variety of fields, we propose four key considerations for visual analysis in display ecologies: 1) Display Composition, 2) Information Coordination/Transfer, 3) Information Connection, and 4) Display Membership. Different aspects of display ecologies stemming from these design considerations will enable users to transform and empower multiple displays as a display ecology for visual analysis.

1 INTRODUCTION

Rapidly advancing technologies have resulted in various types of displays and computing devices for personal and professional use. When a user employs heterogeneous displays collaboratively to achieve a goal, the displays form what is known as a display ecology. Within a display ecology the user can establish the roles of and relationships between displays in a variety of ways. Ultimately, 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, multiple displays can better assist people in enhancing visual analysis with discretized display space and interaction affordances facilitated by the different displays (Figure 1) [3], [6], [7], [8], [10]. Despite its many advantages, 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. Since information and tasks are scattered and disconnected among separate displays, the inherent challenges associated with visual analysis include:

- The user must selectively choose and combine available displays into a holistic visual analysis space to meet the demands of a given analysis task.
- The user must transfer data and visualizations for coordinating information and analysis tasks across displays.
- The user must then connect scattered information across separate displays.
- Throughout an analysis session, the user is likely to dynamically change (add/remove) memberships of available displays in a display ecology to meet the needs of the analysis goal.

Now that analysts are afforded increasing opportunities to conduct their analysis tasks with various displays and devices,

*hchung@cs.uah.edu; †[north, sarang87]@vt.edu; ‡jjchen@umbc.edu



Fig. 1. Display Ecologies for Visual Analysis: Personal Display Ecology (Left) [1] and Collaborative Display Ecology (Right) [3].

these challenges become more important and must be taken into account when developing visual analysis tools based on multiple displays. However, we argue that little research has been undertaken to address these challenges.

The goal of this paper, therefore, is to survey visual analysis techniques that can unite multiple, separate displays and devices into a cohesive analysis space. Toward this goal, we first explored design considerations of display ecologies for visual analysis by surveying and distilling results from prior research in visual analytics, information visualization, sensemaking, distributed cognition, and human-computer interaction. This paper will also help visualization developers and researchers in designing, understanding and evaluating new visual analysis tools that maximize the potential of the numerous displays and devices.

2 METHOD

In order to clarify the design considerations and challenges associated with display ecologies, we conducted a systematic review comprised of the following four steps.

Step 1—Compiling a list of important terms: We began our research with the question: “What interaction and visualization techniques are needed for visual analysis when users are using a display ecology?” To assemble a set of relevant topics associated with interaction and visualization techniques in displays ecologies, we created a list of terms based on the relevant literature addressing both visual analysis and display ecologies. The list included a number of essential characteristics of display ecologies, such as multiple display form factors, physicality, spatiality, input redirection, distributed user interface, multiple views, cross-display object movement, etc. Although display ecologies for visual analysis share some of the challenges (e.g., coordinating and connecting views) in common with taxonomies of multiple views [16], these characteristics implies that visual analysis tools for a display ecology differ in significant ways from the multiple-view visualization.

Step 2—First coding pass: The principal research question and initial set of terms helped us to define our area of focus in reviewing the available literature and in identifying relevant papers. Using these terms, we collected a total of 83 papers from several venues (e.g., VAST, InfoVis, CHI, UIST, ITS, etc.) that specifically focused on visualization and interactive techniques for multiple displays and devices. It should be noted that these papers originated from a variety of communities with a number of different target applications, such as interactive media,

visualization, data analytics, sensemaking, educational platforms, etc. Thus, these systems evidence a fairly diverse set of characteristics. We focused on understanding how users create a visual analysis space in which they can manipulate, coordinate and synthesize data through multiple displays in order to facilitate insight formation. During the first coding pass, we culled several papers that did not pertain to visualization/interaction techniques, systems and user studies for display ecologies. As a result, 61 papers were targeted for further consideration.

Step 3—Open coding: To refine and group more salient dimensions of the design considerations from the 61 relevant papers, we employed Affinity Diagram. During this process, we employed sticky notes to record and spatially organize salient features and ideas pertaining to visual analysis activities (involving, at most, a few words). As we examined and dissected more literature reports via this method, an increasing number of common themes pertaining to 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 challenges and tasks. Ultimately, we identified a set of key dimensions regarding essential design considerations for display ecologies.

Step 4—Validation: The final step involved validating these key dimensions by discussing whether each factor was truly important and relevant. This consensual process helped to narrow the list to what we present herein as the four design considerations of display ecologies.

3 DISPLAY ECOLOGIES FOR VISUAL ANALYSIS

In this section, we define the key ecology concepts informing our design considerations for display ecologies and describe how display ecologies can address visual analytics challenges.

3.1 Defining Display Ecologies

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 that interact with non-living components in the environments, interrelated as a system. As research expands across the disciplines, the notion of an ecology has migrated to other areas. 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 interact with technologies, guided by their own goals and values [20], [22].

We further extend the notion of an ecology by applying it to the modern multi-display environment. When users employ available displays collaboratively to achieve a goal, they mix and match displays to accomplish their desired outcomes. A group of such displays forms 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, we define a *Display Ecology* as a system of displays that engage the entire workflow of a task to better assist analysts in achieving their desired outcomes.

A diversity of displays in terms of their size, mobility, and inherent affordances all play important roles in forming a display ecology. Small mobile displays can accompany people as they move about, while larger displays inhabit specific spaces and can form specific ecosystems. Different display form factors such as wall displays or tabletop surfaces support different types of analysis activities and multiple collaborating users. In this regard,

prior research has offered guidelines about how different types of displays and devices (e.g., large high resolution displays and tabletops) support visual analysis tasks [25], [27]. However, rather than using discrete displays, recent studies in the domain of HCI research are largely concerned with how an ensemble of different displays can better assist people and their tasks in fulfilling various applications [6], [30], [33]. These studies confirm that the capabilities of each display to support tasks can be fully realized when multiple displays function in concert with each other.

3.2 Addressing Challenges of Visual Analytics

In this section, we present five significant challenges associated with visual analytics. Although these challenges have been partially identified and discussed in prior visual analytics studies [34], few of them have adequately addressed how these issues can be effectively managed when utilizing multiple displays. We therefore discuss these challenges in the context of display ecologies.

Scalability: Scalability issues represent one of the biggest visual analytics challenges [34]. While the integration of visualization with computational methods will aid analysts to illuminate patterns, trends, clusters, outliers, etc. in their data, the amount of information that can be simultaneously visualized is still hindered by the physical constraints of the display medium.

Similar to large tiled-wall displays, display ecologies provide a larger collective display space, but have an inherent advantage in that they can facilitate better utilization of physical space beyond one single virtual raster space. Moreover, the various displays can be located at different angles, as well as in different places. Prior research has shown that the use of physical spaces impacts insight formation significantly [8], [35].

Heterogeneous Data: An increasing number of visual analytics tasks must facilitate seeing, comparing and interpreting heterogeneous data from different sources in order to solve analysis problems (e.g., in an intelligence analysis setting). Enabling analysts to understand information derived from heterogeneous sources remains a visual analytics challenge [34].

Display ecologies can support the presentation of heterogeneous information for visual analysis, since multiple displays enable analysts to see, combine, and interact with diverse and often complex types of information via separated displays with increased screen space and different interaction techniques [19].

Analytic Provenance: Another challenge in the analytic process lies in keeping track of analysis tasks and data changes that users have made during an analysis sessions. It is important for users to understand how their analytical steps were formed and ultimately derived a final result [40]—not only to validate results, but also to be able to reproduce their results.

In a display ecology, each component display allows the user to distribute different data, workspaces, and analysis tasks into independent displays. The multiple, separate displays facilitate the continued visual availability of information and the maintenance of analysis tasks, helping users to understand where data has come from and how it has contributed to the entire result [39].

Collaborative Analysis: Collaborative analysis results often represent a combination of both individual and collaborative efforts. Even though users may be working on collaborative tasks, their individualized tasks are essential to a collective outcome. However, synthesizing the analytic results of both individual and collaborative tasks can be particularly challenging, often requiring significant effort to manage information from individual, disparate analyses [3], [7].

The separate displays facilitate the ability for collaborating users to distribute individual and collaborative work effectively with individual and shared displays [3], [29].

Table 1. Four design considerations and their relevant analysis tasks and challenges.

Design Considerations	Sub-category	Related Visual Analysis Tasks	Addressing Challenges
Display Composition	Distributed Controls	<ul style="list-style-type: none"> Remote Selection/Filter/Navigation: Remotely isolate and interact with a subset of data items using UIs and input devices on different displays [2], [4]. Search: Find related information allowing the users to explore the data independently from others [5], [4], [2]. 	Collaborative Analysis Opportunistic Analysis
	Data View: Single Continuous	<ul style="list-style-type: none"> Physical Navigation: Physically navigate the increased visualization space from multiple displays [9], [11], [12], [13]. 	Scalability
	Data View: Multiple Coordinated	<ul style="list-style-type: none"> Compare: Juxtapose visualizations and data items [14], [15]. Complement: See different aspects of a dataset [17], [18]. Split: Decompose a visualization view / entity into multiple views on multiple displays [19]. 	Heterogeneous Data
	Data View: Navigation Metaphor	<ul style="list-style-type: none"> Focus + Context: Multiple displays are combined to represent Focus+Context metaphors [21], [23], [24]. Overview + Detail [26], [28], [29]. 	Heterogeneous Data Collaborative Analysis
Information Coordination/Transfer	Data View Semantic Substrates	<ul style="list-style-type: none"> Organize/Space-to-think [31]: Spatially organize data items and views across displays [32], [3], [8]. 	Sensemaking Analytic Provenance
	Synchronize Surrogate Nominal Physical	<ul style="list-style-type: none"> Publish: Broadcast and synchronize updates among different visualizations [28], [19]. Share and Transfer: Share data items and meta content such as annotations, bookmarks to other displays [36], [37], [38]. 	Sensemaking Heterogeneous Data
Information Connection	Overview	<ul style="list-style-type: none"> Merge: Automatically aggregate multiple visualizations from different displays and visual analysis results [39], [18], [41]. 	Analytic Provenance
	Explicit	<ul style="list-style-type: none"> Brushing and Linking: Select and see related visual entities across displays [8], [1], [17], [19]. Annotate / Bookmark: Guide important information on different displays and devices/Add metadata to visualization elements [1]. 	Heterogeneous Data
	Implicit	<ul style="list-style-type: none"> Organize/Space-to-think [31]: Create Semantic Structures across displays [32], [3], [8]. 	Sensemaking
Display Membership	Pre-designed	<ul style="list-style-type: none"> Divide: Design the division of specific analysis tasks and information based on more suitable displays and users [29]. 	Collaborative Analysis
	Ad-hoc	<ul style="list-style-type: none"> Branch: Branch visual analysis tasks on different displays to explore and analyze the data on other available displays anytime and anywhere [44], [13], [14], [38], [18], [28]. 	Opportunistic Analysis

Opportunistic Analysis: While conducting visual analytics tasks, analysts may encounter, retain, and consult various pieces of information opportunistically across diverse spaces throughout the day and week (or longer). Insight formation may emerge from the domain knowledge of the analyst by chance; it can be built up over time and is often unpredictable and serendipitous [42].

Multiple displays can accommodate an analyst’s changing analysis needs by enabling the user to combine or shift different displays and devices, thereby tapping into the specific potential of different types of technologies as needed [43].

Sensemaking: Visual analytics emphasizes more cognitively demanding issues [34]. It can provide computational methods for grappling with large-scale document data, such as filtering, clustering, identifying potentially important information and making connections. Nonetheless, sensemaking is still crucial because such strategies do not necessarily result in full cognition—although they certainly contribute to an analyst’s understanding of data.

Sensemaking tasks can be extended to display ecologies with multiple screen spaces and cross-device interactions. Utilizing multiple displays enable users to further externalize their cognitive reasoning processes. In other words, users can leverage multiple displays to spatially organize and structure information for their sensemaking process [3], [8].

4 FOUR DESIGN CONSIDERATIONS

In this section we explore each of the four design considerations by which users perform different analysis tasks using multiple displays. Based on the characteristics of visual analysis activities from our prior studies and available visual analysis research, we can segregate analysis tasks with multiple displays into four main considerations:

- **Display Composition:** defines relationships among combined displays to create an integrated visual analysis space.
- **Information Coordination/Transfer:** corresponds to how analysis tasks, visualization, and data can be moved and distributed among displays.
- **Information Connection:** corresponds to identifying a design consideration to enable users to connect and integrate scattered information across different displays.

- **Display Membership:** captures the differences between devices in a display ecology according to prescribed design, as well as the addition or removal of different displays during analysis.

There are two primary reasons why we target these four factors in this article. First, we believe that several visual analytics challenges (§3.2) can be addressed via our design considerations. Specifically, each design consideration explores the important challenges and tasks for visual analysis. Table 1 lists the relationships between our design considerations and relevant visual analysis tasks/challenges.

Second, our design considerations involve the sensemaking and cognitive reasoning process by which information is iteratively collected, organized, presented and synthesized in order to form insight and knowledge [45]. In general, such processes consist of different phases, each of which requires different analysis and interaction techniques. The basic structure of such analysis activities within a display ecology can be illustrated by the following scenario.

An analyst wants to analyze and understand a large weather dataset with multiple displays. She first opens a multiple-view visualization (e.g., maps, atmospheric pressure, temperature changes, etc.) on a single display; however, the resolution and size of a single display may not be sufficient to show details of the multi-view visualization. Thus, she needs to enlarge each view of the multi-view visualization and decides to *choose and combine* three available displays at her office. She then *coordinates* the three visualizations onto the displays: a desktop screen, projector, and laptop based on the data size. After coordinating the multi-view visualization with different displays, she needs to understand how visual items on each display are related to the others on the additional screens. She accomplishes this by highlighting the information across displays. Finally, she can *synthesize* all of the connected information from displays in order to broaden her understanding of otherwise 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.

4.1 Display Composition

The first step in performing analysis with 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. We classify such inter-display relationships into the following two dimensions, based on how users form relationships among displays for visual analysis.

4.1.1 Distributed Controls

A multi-display relationship can be considered to incorporate “distributed controls” when a display ecology allows users to distribute the analysis tool elements—including the primary visualization view and user interfaces—into different displays, either 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 [5]. The user’s main analysis is performed with a specific primary display, while the other displays play a supporting role for the primary display in terms of their interactions and data. Generally, primary displays (e.g., a large display) are employed to visually explore and analyze data as a visualization view, while the supporting displays (e.g., one or more mobile display) serve as remote controllers [2], [4]. A major advantage of this approach is that it facilitates physical navigation and collaborative inputs, such that users can independently control visualizations. However, the distribution of controls and elements into different displays leads users to repeatedly switch displays and contexts between the supporting and the primary displays.



4.1.2 Data Views

We can define a display ecology in terms of composite visualization views or structures. Thus, the data view relationship focuses on how integrated visualization views and structures can be constructed with multiple displays. We present four different examples of this relationship for constructing data views with separate displays.


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 for a single visualization. Examples of this relationship include the dynamic tiling display [9], the pass-them-around prototype [11], and the peer-to-peer distributed user interfaces [12]. These approaches enable the user to combine random mobile screens into one single tiled display.




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 [19]. 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 can be easily 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 that users can better exploit the multiple-display space to see different aspects of data. However, optimizing each visualization view for different displays is challenging and it can be problematic to implement brushing and linking across devices running different computing platforms.



Navigation Metaphor: A display relationship can be considered to be in “navigation metaphor” when more than two displays are combined to represent InfoVis 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 data. For example, by using spatially-aware mobile devices in conjunction with large shared displays, users are able to form a display ecology for enhanced navigation views such as Tangible Views [21], and iPodLoupe [23]. It should be noted that this approach requires special devices to track and detect user gestures and physical actions (e.g., motion trackers and touch devices).



Semantic Substrate: A multiple-display relationship can be considered to be a “semantic substrate” when each display in an ecology maintains a spatially separate, non-overlapping screen space in which different information or visual items are laid out based on related topics. Generally, this relationship does not emphasize the creation of 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 [3], [32], [8]. Displays in such a relationship enable users to separate data semantically across displays, as well as to distribute analysis tasks. However, because analysis tasks are performed within each display separately, the complexity of synthesizing distributed information is increased.



4.2 Information Coordination/Transfer

The multiple-screen space facilitated by display ecologies inherently allows one or more users to divide analysis tasks into small parts based on display properties. To divide and coordinate 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.

4.2.1 Synchronized

A coordination technique can be considered to be “synchronized” when information coordination is accomplished through synchronizing interactions and data on each display through input redirections or a networked database. This approach automatically propagates any visual analysis operation performed on one display to all the other displays. Therefore, information can be shared among displays without additional interaction. With this approach, users may perform work via shared focus of the visualization view and simultaneous individual control of the dataset on each device [19], [28], [3]. A typical visual analysis task that would benefit from this approach is when each individual on a team is carrying out his or her data analysis from one display—but needs to be aware of how other team members are manipulating the data. However, the introduction of unnecessary “noise” associated with an increasing number of displays is one of the chief drawbacks of this approach in collaborative analysis environments, since displays in the display ecology may share everything instantly and indiscriminately.



4.2.2 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 example approaches for surrogacy are via a virtual lens or a display proxy.



Virtual Windows: Coordinating information can be accomplished using a large shared virtual space. Each display is represented by a localized window (which represents the view of that particular display), which is projected onto the shared virtual space. For instance, the Zoomable Object-oriented Information Landscape (ZOIL) [36] is a multi-display 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. Thus, a user can coordinate and show visual items on different displays by indirectly moving the associated lenses in the shared virtual space.

Display Proxy represents a virtual destination display for transferring items, as well as facilitates the availability/connectivity of different displays in an ecology. It provides a virtual reference for one or more displays on the main display. For example, Wigdor et al. designed an interaction space that uses multiple wall-mounted displays and one tabletop display [37]. Each wall-mounted display features the World in Miniature (WIM) view, which is a corresponding proxy for the different displays. If a user wants to relocate one item to a different display, he or she can drag it to the corresponding proxy on a display. A drawback with this approach, however, is that visual items can become too small to be legible and manipulated on a proxy, requiring compensation through zooming and/or other interaction techniques.

4.2.3 Nominal

A nominal reference (e.g., a document ID and file name) of both data and displays can be employed in order to coordinate and transfer information or tasks among displays. Using nominal techniques, information sharing requires users to memorize nominal information such URL, filename or display IDs. When users want to send or move specific data 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 content. This approach represents the predominant way for sharing files and information with others in current desktop environments. Multibrowsing [46] and Conductor [8] are examples of this approach. The approach forces users to know and remember the names of the desired displays. With this approach, coordinating information and data often requires going through one or more indirect and supplementary procedures (e.g., checking display names to send information).

4.2.4 Physical

The goal of the “physical” coordination technique is to enable the user to coordinate data items across displays based on the physical reference of each display using gestural techniques. This approach emphasizes 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”, or “tilting” and “throwing.” 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 require users to directly place one device in contact with another device to transfer information between the two devices [47].




A primary benefit of this approach is that when a user needs to coordinate a specific data item, she or he can focus on its content. Notions that may be important in providing such an experience include making the data user-focused rather than device-focused and tailoring devices to their roles in display ecologies. However, because this approach depends on a user's physical actions, it runs the risk of inaccuracy in that a user may inadvertently send specific information to an unintended display.

4.3 Information Connection

By design, the use of a display ecology means that information of interest and analytical activities are distributed over different displays [33]. The principal challenge associated with visual analysis in a display ecology is 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. In short, a display ecology user must connect and integrate relevant information across displays in order to synthesize diverse sources of data. In this section, we explore the design consideration associated with how to connect and integrate information from different information and data items (often involving different visual representations) over separate displays.

4.3.1 Overviews

A connection 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 display. This approach is designed to directly integrate visualization components and views from different displays in an ecology onto a single display.



Each user's analysis activities contribute to creating an overview visualization, which facilitates an awareness of other's progress or the connections among displays within a single integrated view. It also helps to compare and visually connect data content by showing all related information side-by-side on a single screen [41]. Generally, users do not actively create this visualization; rather, the system automatically creates the visualization using the information that users generate. For example, IdeaVis [39] 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.

On the other hand, information and visual items from different displays can be manually overlaid or merged on a single display. For instance, WeSpace allows users to organize and overlay images created on different displays onto a large display simultaneously [18]. With VisPorter, users can merge visualization components created from different displays to an integrated visualization on a single large display [3].

One challenge inherent in this approach is related to how the visual analysis system can appropriately capture users' analysis activities and data changes from different displays in a single overview screen.

4.3.2 Explicit Connections

Information on separate displays can be visually connected with different representations (line, shapes, color, etc.) spanning multiple displays. The visual representations denote a method for elucidating the visual connections between the data items as a single visual structure. The challenge is how a visual representation might be extended in a display ecology, whereby analysts can be directed to relevant information across displays.



We now consider two visual representations for accomplishing this goal.

Highlighting: The simplest way to represent data connections across displays is to highlight one data point or a set of data points or visual items with or without some labelled text. These highlights enable users to discriminate linked data items located on different displays with color or shapes. The highlighting connects data items that represent the same information across displays, thereby clarifying how those items are related in the different visualizations or analysis workspaces from different displays. 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.

Cross-Display Visual Representations: The cross-display visual representations are grounded in the “partially out of the frame” approach presented by several off-screen visualizations such as Halo [48]. A simple but effective form of cross-display visual representation is the visual link, which corresponds to one or more straight lines from a source to multiple targets across displays [1]. Specifically, a cross-display link exits or enters through one of the off-window directions, while a portion of that link protrudes into another display’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, they can give the illusion of one continuous workspace utilizing discrete displays. The advantage of employing one or more visual representations across displays is to enable users to connect information artifacts from different displays, while at the same time maintaining their analysis context on each display without switching displays. This approach, however, tends to result in many visual objects on each screen, particularly if there are many items related across displays.

4.3.3 Implicit Connections

The user can also externalize and clarify relationships of scattered information by organizing it through spatial relations among data, as well as displays. The implicit technique enables the analyst to leverage the multiple displays for externalization of cognitive data synthesis. The multiple discretized screen space allows the user to extend the concept of “space to think” [31]. The user can spatially organize data among displays to construct semantic structures for scattered information, thereby transforming the random layout of data items on each display into semantically-meaningful structures (e.g., regions, people, timelines, etc.). Several recent studies have investigated how users form new knowledge about a complex set of data using this approach, such as large document datasets, and in support of the sensemaking process [8], [36], [3].

With this approach, the analyst cognitively connects and integrates diverse pieces of relevant information from different displays based on their cross-display semantic structure or clusters. However, the challenge associated with this technique is to maintain awareness of information from different data sources over separate displays.

4.4 Display Membership

The pre-designed and ad-hoc display ecologies can generally be distinguished by the dynamic membership of each display, as characterized by the number of displays, types of displays, and task division.

4.4.1 Pre-Designed Display Ecologies

A display configuration can be considered to be 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 specific analytical tasks. Thus, the role or tasks of different displays are inherently fixed by design.

The main goal of the 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 leverage better specific display properties and settings for analysis tasks. For example, a user will be able to forage for information only on his or her personal displays—but then multiple users can merge their information collaboratively on a large display. An illustration of this scenario is the Pixel-oriented Treemap for multiple display [29], which is designed to divide two different visualization tasks between two types of displays (overview and detail).

While this approach enables users to better exploit the specific visual and analysis capabilities of different types of displays, it must be stressed that users cannot dynamically reorganize the analysis space to accommodate frequently-changing needs or available displays.

4.4.2 Ad-hoc Display Ecologies

As noted earlier, the growing availability and complexity of both devices and data—coupled with the urgency of certain analysis tasks—means that analysts will be called upon increasingly to engage with diverse pieces of information and devices at opportunistic moments (e.g., the identification of terrorist plots or analysis of an outbreak of infectious disease). Such scenarios will inevitably require the formation of an “ad-hoc display ecology” that emphasizes the dynamic reorganization of available displays for different analysis requirements. In this approach, a display ecology can be formed with available heterogeneous displays opportunistically. In contrast to the pre-designed configuration, this ecology focuses on creating analysis space by dynamically assigning different tasks to and combining available displays as analysis requirements evolve. In this way, the user can deploy and span analytic tasks across different types of available displays in adaptable configurations. Since analysis is not confined to a specific display, the analysis space can consist of various types of additional displays, including 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 [44], [43].



Building visual analysis tools based on multiple heterogeneous devices and ad-hoc ecologies is difficult due to system-imposed constraints, e.g. the heterogeneity of communication protocols and different software and hardware platforms. Therefore, it becomes essential for a software infrastructure to support this interoperability among different displays and devices.

5 DISCUSSION

In this section, we discuss how our four considerations can further augment the design of visual analysis tools in a display ecology.

5.1 Balance Foraging and Synthesis Approaches

While some of the approaches in our considerations are useful for navigating and exploring data, others focus more on facilitating analysts’ cognitive reasoning and sensemaking processes. We can further divide our design considerations along a spectrum of foraging-oriented and synthesis-oriented approaches. 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 are related to sub-dimensions of the data views (§4.1.2), such as the single continuous view and navigation metaphor. This approach works best when the goal is to spend a considerable portion of an analysis searching, filtering, reading, collecting, and visualizing information using multiple displays. We recommend that this approach be used to increase the overall screen real estate in order to see more data in geographical and multiple-view visualization applications enabled by display ecologies.

On the other hand, a synthesis-oriented approach concentrates on cognitive issues associated with the sensemaking process of deciphering information. Specifically, a synthesis-oriented approach emphasizes externalizing the user's thought processes by organizing and distributing the collected information on multiple displays. For example, by utilizing both the semantic substrate (§4.1.2) and implicit connections (§4.3.3), 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. [31] and Robinson [49] confirmed, spatially arranging documents into increasingly formal and meaningful structures (i.e., spatial clustering or ordering) enables one to externalize the sensemaking process.

In general, visual analysis processes comprise a broad range of analysis activities, including foraging and synthesis phases. However, little is known about how these two approaches can be balanced in a display ecology towards the goal of promoting efficiency in exploring and synthesizing large datasets. Hence, one important research avenue for visual analysis in a display ecology would be to investigate how to balance those two approaches to support the entire analysis process.

5.2 Exploit the Physicality of Display Ecologies

Our physical space is crucial for insight formation since we are embodied beings who live in the physical world [50]. In display ecologies, the physical properties of displays (e.g., physical shape, size, specific form factors, etc.) and their spatial properties (e.g., location, orientation, proximity, topology of displays, etc.) guide users toward which inter-display interactions are possible. Several approaches from our design considerations leverage both the physicality and spatiality afforded by a display ecology to support visual analysis tasks.

For example, the navigation metaphor approaches (§4.1.2) employ physical affordances of mobile devices and tabletop surfaces by placing the phone or tablet directly in contact with a large tabletop in order to navigate and transfer data [47]. The mobile displays are spatially relocated to customize and adapt visualization views for data exploration tasks. Also, the physical techniques for the information transfer (§4.2.4) rely on the physical reference and presence of displays to move and coordinate visual items in a display ecology [51]. For these tools and techniques, the physicality of each display plays a crucial role in how to construct an analysis workspace (§4.1) as well as how to achieve the desired interactions between displays.

The physicality of display ecologies facilitates visual analysis by engaging and employing the user's physical skills. In this regard, further research should be conducted about how to better tap into the physical and spatial characteristics of multiple displays.

5.3 Additional Considerations

It must be noted that there may be other, more specialized, visual analysis activities that need to be incorporated—either for performance reasons or for unique analysis scenarios. Thus, the design considerations discussed herein should not be considered to be all inclusive. Throughout this article, we have acknowledged

that the four considerations presented herein can be improved by discussing new aspects of display ecologies. For instance, nowadays there are many types of available devices for both home and office—ranging from very large display screens (e.g., touch table) to extremely small displays (e.g., wearable devices). It is worth discussing which types of displays can better serve a specific visual analysis task or approach. Thus, classifying the specific displays can aid in identifying more appropriate techniques for different analysis needs.

Most of the approaches discussed in this article focus on display ecologies constructed by user's needs and manual interactions. However, we believe that further research needs to be conducted about how to support and balance between automatic support and manual adjustments to coordinate, transfer, and connect information across displays for visual analysis.

6 CONCLUSION

Our current computing environment can benefit from leveraging a large number of available displays to explore and analyze large and complex data aggregates. In this paper, we explored essential considerations of display ecologies that enable users to transform and empower multiple displays as a display ecology for a cohesive visual analysis space. When considered collectively, these considerations and the related techniques can enhance our understanding of how to accomplish visual analysis in a display ecology.

In particular, one of the notable benefits of this investigation is that the four considerations make it possible for a display ecology to offer improved visual analysis environments and tools—but with a significantly reduced price-tag, since users will be able to combine readily accessible displays in their workspace.

We believe that this article presents critical components for analyzing a large amount of information and data via the use of a display ecology. We hope that our survey and discussions will help guide the creation of new visual analysis systems for display ecologies and provide inspiration for future research in visual analysis.

ACKNOWLEDGEMENT

This work was partially supported by NSF IIS-1218346, a grant from Department of Defense, NSF DBI-1260795, IIS-1302755, MRI-1531491, NIST MSE-20NANB12H181, and DoD USAMRAA-13318046. We thank Adam Barnes for illustrating Figures.

REFERENCES

- [1] H. Chung, "Designing Display Ecologies for Visual Analysis," Ph.D. Dissertation, Dept. of Computer Science, Virginia Polytechnic Institute and State University, Blacksburg, VA, 2015.
- [2] Y. Jansen, P. Dragicevic, and J.-D. Fekete, "Tangible Remote Controllers for Wall-Size Displays," *Proc. ACM CHI'12*, pp. 2865-2874, 2012.
- [3] H. Chung, C. North, J. Z. Self, S. Chu, and F. Quek, "VisPorter: facilitating information sharing for collaborative sensemaking on multiple displays," *Personal and Ubiquitous Computing*, vol. 18, pp. 1169-1186, 2014.
- [4] O. Chapuis, A. Bezerianos, and S. Frantziskakis, "Smarties: An input system for wall display development," *Proc. ACM CHI'14*, pp. 2763-2772, 2014.
- [5] J. Yang and D. Wigdor, "Panelrama: enabling easy specification of cross-device web applications," *Proc. ACM CHI'14*, pp. 2783-2792, 2014.
- [6] E. M. Huang, E. D. Mynatt, and J. P. Trimble, "Displays in the wild: understanding the dynamics and evolution of a display ecology," presented at *Pervasive Computing'06*, 2006.

- [7] P. Isenberg, T. Isenberg, T. Hesselmann, B. Lee, U. Von Zadow, and A. Tang, "Data Visualization on Interactive Surfaces: A Research Agenda," *IEEE CG&A*, vol. 33, pp.16–24, 2013.
- [8] P. Hamilton and D. J. Wigdor, "Conductor: enabling and understanding cross-device interaction," *Proc. ACM CHI'14*, pp. 2773-2782, 2014.
- [9] M. Li and L. Kobbelt, "Dynamic tiling display: building an interactive display surface using multiple mobile devices," *Proc. Conf. on Mobile and Ubiquitous Multimedia*, pp. 24, 2012.
- [10] A. Radloff, C. Tominski, T. Nocke, and H. Schumann, "Supporting presentation and discussion of visualization results in smart meeting rooms," *The Visual Computer*, pp. 1-16, 2014.
- [11] A. Lucero, J. Holopainen, and T. Jokela, "Pass-them-around: collaborative use of mobile phones for photo sharing," *Proc. ACM CHI'11*, pp. 1787-1796, 2011.
- [12] E. R. Fisher, S. K. Badam, and N. Elmqvist, "Designing peer-to-peer distributed user interfaces: Case studies on building distributed applications," *IJHCS*, vol. 72, pp. 100-110, 2014.
- [13] 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," *Proc. ACM ITS'12*, pp. 235-244, 2012.
- [14] S. K. Badam, E. Fisher, and N. Elmqvist, "Munin: A Peer-to-Peer Middleware for Ubiquitous Analytics and Visualization Spaces," *IEEE TVCG*, vol. 21, pp. 215-228, 2015.
- [15] M. Beaudouin-Lafon, S. Huot, M. Nancel, W. Mackay, E. Pietriga, R. Primet, et al., "Multisurface Interaction in the WILD Room," *Computer*, vol. 45, pp. 48-56, 2012.
- [16] M. Q. W. Baldonado, A. Woodruff, and A. Kuchinsky, "Guidelines for using multiple views in information visualization," *Proc. AVI'00*, pp. 110-119, 2000.
- [17] 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," *Proc. ACM CHI'09*, pp. 1237-1246, 2009.
- [18] 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," *Proc. ACM CHI'08*, pp. 1357-1360, 2008.
- [19] S. K. Badam and N. Elmqvist, "PolyChrome: A Cross-Device Framework for Collaborative Web Visualization," *Proc. ACM ITS'14*, pp. 109-118, 2014.
- [20] K. Krippendorff, *The semantic turn: A new foundation for design*: CRC Press, 2005.
- [21] M. Spindler, C. Tominski, H. Schumann, and R. Dachselt, "Tangible views for information visualization," *Proc. ACM ITS'10*, pp. 157-166, 2010.
- [22] B. A. Nardi and V. O'Day, *Information Ecologies: Using Technology with Heart*: MIT Press, 1999.
- [23] S. Vaida, 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," *Proc. ACM ITS'09*, pp. 109-116, 2009.
- [24] J. C. Lee, S. E. Hudson, J. W. Summet, and P. H. Dietz, "Moveable interactive projected displays using projector based tracking," *Proc. UIST'05*, pp. 63-72, 2005.
- [25] C. Andrews, A. Endert, B. Yost, and C. North, "Information visualization on large, high-resolution displays: Issues, challenges, and opportunities," *Information Visualization*, vol. 10, pp. 341-355, 2011.
- [26] C. Forlines, A. Esenther, C. Shen, D. Wigdor, and K. Ryall, "Multi-user, multi-display interaction with a single-user, single-display geospatial application," *Proc. ACM UIST'06*, pp. 273-276, 2006.
- [27] P. Isenberg, U. Hinrichs, M. Hancock, and S. Carpendale, "Digital tables for collaborative information exploration," in *Tabletops-Horizontal Interactive Displays*, ed: Springer, pp. 387-405, 2010.
- [28] J. Sanneblad and L. E. Holmquist, "Ubiquitous graphics: combining hand-held and wall-size displays to interact with large images," *Proc. AVI'06*, pp. 373-377, 2006.
- [29] H. Chung, Y. J. Cho, J. Self, and C. North, "Pixel-oriented Treemap for multiple displays," *Proc. IEEE VAST'12*, pp. 289-290, 2012.
- [30] 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," *IJHCS*, vol. 70(10), pp. 765-779, 2012.
- [31] C. Andrews, A. Endert, and C. North, "Space to think: large high-resolution displays for sensemaking," *Proc. ACM CHI'10*, pp. 55-64, 2010.
- [32] F. Geyer and H. Reiterer, "A cross-device spatial workspace supporting artifact-mediated collaboration in interaction design," *Proc. ACM CHI '10 Extended Abstracts*, pp. 3787-3792, 2010.
- [33] D. Dearman and J. S. Pierce, "It's on my other computer!: computing with multiple devices," *Proc. ACM CHI'08*, pp. 767-776, 2008.
- [34] J. J. Thomas and K. A. Cook, *Illuminating the Path: The Research and Development Agenda for Visual Analytics*: National Visualization and Analytics Center, 2005.
- [35] R. Ball, C. North, and D. A. Bowman, "Move to improve: promoting physical navigation to increase user performance with large displays," *Proc. ACM CHI'07*, 2007.
- [36] H.-C. Jetter, M. Zöllner, J. Gerken, and H. Reiterer, "Design and Implementation of Post-WIMP Distributed User Interfaces with ZOIL," *IJHCI*, vol. 28, pp. 737-747, 2012.
- [37] D. Wigdor, C. Shen, C. Forlines, and R. Balakrishnan, "Table-centric interactive spaces for real-time collaboration," *Proc. AVI'06*, pp. 103-107, 2006.
- [38] N. Marquardt, K. Hinckley, and S. Greenberg, "Cross-device interaction via micro-mobility and f-formations," *Proc. ACM UIST'12*, pp. 13-22, 2012.
- [39] F. Geyer, J. Budzinski, and H. Reiterer, "IdeaVis: a hybrid workspace and interactive visualization for paper-based collaborative sketching sessions," *Proc. NordiCHI'12*, pp. 331-340, 2012.
- [40] C. North, R. Chang, A. Endert, W. Dou, R. May, B. Pike, et al., "Analytic provenance: process+ interaction+ insight," *Proc. ACM CHI'11 Extended Abstract*, pp. 33-36, 2011.
- [41] C. Eichner, T. Nocke, H.-J. Schulz, and H. Schumann, "Interactive Presentation of Geo-Spatial Climate Data in Multi-Display Environments," *ISPRS International Journal of Geo-Information*, vol. 4, pp. 493-514, 2015.
- [42] R. Chang, C. Ziemkiewicz, T. M. Green, and W. Ribarsky, "Defining Insight for Visual Analytics," *IEEE CG&A*, vol. 29, pp. 14-17, 2009.
- [43] N. Elmqvist and P. Irani, "Ubiquitous analytics: Interacting with big data anywhere, anytime," *Computer*, pp. 86-89, 2013.
- [44] R. Rädle, H.-C. Jetter, N. Marquardt, H. Reiterer, and Y. Rogers, "HuddleLamp: Spatially-Aware Mobile Displays for Ad-hoc Around-the-Table Collaboration," *Proc. ACM ITS'14*, pp. 45-54, 2014.
- [45] 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.
- [46] B. Johanson, S. Ponnekanti, C. Sengupta, and A. Fox, "Multibrowsing: Moving Web Content across Multiple Displays," *Proc. Ubicomp'01*, pp. 346-353, 2001.
- [47] D. Schmidt, J. Seifert, E. Rukzio, and H. Gellersen, "A cross-device interaction style for mobiles and surfaces," *Proc. ACM DIS'12*, pp. 318-327, 2012.
- [48] P. Baudisch and R. Rosenholtz, "Halo: a technique for visualizing off-screen objects," *Proc. ACM CHI'03*, pp. 481-488, 2003.
- [49] A. C. Robinson, "Collaborative synthesis of visual analytic results," *Proc. IEEE VAST'08*, pp. 67-74, 2008.
- [50] W. Kuhn, "Handling data spatially: Spatializing user interfaces," *Proc. 7th International Symposium on Spatial Data Handling (SDH'96)*, pp. 877-893, 1996.
- [51] 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.