Pers Ubiquit Comput (2018) 22:409–431
https://doi.org/10.1007/s00779-017-1091-4

ORIGINAL ARTICLE

SAViL: cross-display visual links for sensemaking in display ecologies

Haeyong Chung[1] · **Chris North**[2]

Received: 29 January 2017 / Accepted: 27 October 2017 / Published online: 12 December 2017

Abstract The main challenge associated with visual analysis using multiple displays is tied to the fact that a user must maintain awareness of and synthesize scattered information across separate displays—some of which may be out of the user's immediate field of vision. To address this need, we present Spatially Aware Visual Links (SAViL), a cross-display visual link technique capable of (1) guiding the user's attention to relevant information and (2) visually connecting related information across displays. In essence, SAViL visually represents the direct connections among different types of visual objects on separate displays to help users create semantic layers of documents spread over different displays. To test the efficacy of this system, we evaluated the impact of visual linking on the sensemaking process for text data utilizing multiple heterogeneous displays. The results of our evaluation indicate that cross-display links enable users to effectively forage for, organize, and synthesize relevant information scattered across multiple displays, integrating the different displays into a single cohesive visual workspace to support their sensemaking tasks.

Keywords Sensemaking · Multi-display environment · Display ecology · Visual text analytics · Visual links

✉ Haeyong Chung
haeyong.chung@uah.edu

[1] Department of Computer Science, University of Alabama in Huntsville, 301 Sparkman Drive, Technology Hall, Room N341, Huntsville, AL 35805, USA

[2] Department of Computer Science, Virginia Tech, 114 McBryde Hall, Blacksburg 24061, VA, USA

1 Introduction

The current proliferation of various types of displays and mobile devices has significantly expanded their potential for personal, professional, and public use. Increasingly, these heterogeneous displays in our both personal and professional workspaces are functioning in concert to assist us in achieving a desired analysis outcome through the formation of a *display ecology* [1]. A diversity of displays in terms of size, mobility, and inherent affordances all play unique roles in supporting different tasks and goals—especially when used together as a display ecology.

In particular, a display ecology can present a number of exciting opportunities for sensemaking [2, 3]. A range of sensemaking studies involving the use of a display ecology showed that semantically separated space among multiple displays facilitates enhanced understanding and organization of the document data. For instance, the use of a display ecology can facilitate sensemaking tasks and visual analysis, since multiple displays enable users to (1) utilize separate screen space for semantic layers [4–6], (2) tap into the potential of different types of technologies for suitable tasks [7–9], and (3) collaborate more flexibly by enhancing the needs of multiple users through multiple displays [10, 11].

In spite of these benefits, the significant challenge associated with using multiple displays for sensemaking of text data is to maintain awareness of, and subsequently connect/integrate, relevant information over separate displays. Several of these displays may be beyond the user's immediate visual field [1, 12]. However, since information and tasks should be scattered and disconnected among separate displays, an inherent design challenge associated with visual analysis in a display ecology is enabling users to seamlessly coordinate and subsequently connect and synthesize information across displays.

Previously, we began exploring this challenge by building and evaluating *VisPorter*, a multi-display system for sensemaking [6]. As observed in our VisPorter study, when collaborating users performed a sensemaking task with multiple displays, we observed that the use of multiple displays required the user to switch frequently among multiple foci of interest (an average of 33.1 documents were organized and scattered over three different displays). This complicated the ability to mentally connect and integrate scattered information in order to generate a cohesive story. Importantly, the participants pointed out their need for additional features in order to link information across displays revealing relationships; they wanted to verify relationships among findings across the displays during the analysis.

Many of the current visual analysis systems based on 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 [13–15]. Although these highlighting approaches can make it easier for the user to distinguish relevant information from data, the user must rely solely on memory to locate and compare various pieces of information on different displays. This issue becomes more problematic when the amount of information and the number of displays and devices are increased. If workspaces are altered in a display ecology, users may forget the location of pertinent information. Furthermore, the highlighting approaches can discriminate linked data items located on different displays with limited colors, so users can perceive only a small number of connections among these items on multiple displays [16, 17]. Thus, current highlighting techniques are less effective for showing diverse relationships between multiple data elements scattered across more than two displays.

To address these problems, we present Spatially Aware Visual Links (SAViL), which is capable of (1) guiding the user's attention to relevant information across displays and (2) visually connecting related information among separate displays in order to support sensemaking tasks in various display ecologies (Fig. 1). Inspired by the spatial and physical characteristics of multiple displays, SAViL augments simple visual link approaches that connect document elements or visual objects (Fig. 2) of text documents (i.e., keywords, documents, and collections of documents) across displays in a tiled display configuration—providing that the individual displays are proximally located to display more information simultaneously.

The main goal of SAViL is to help the user perform sensemaking tasks by embedding and organizing visual data items onto different displays with simple visual links. It combines the concept of visual links among entities and documents, which have been demonstrated to have positive effects on sensemaking of text documents [16–19]. Importantly, SAViL focuses on supporting *Space to Think* for sensemaking

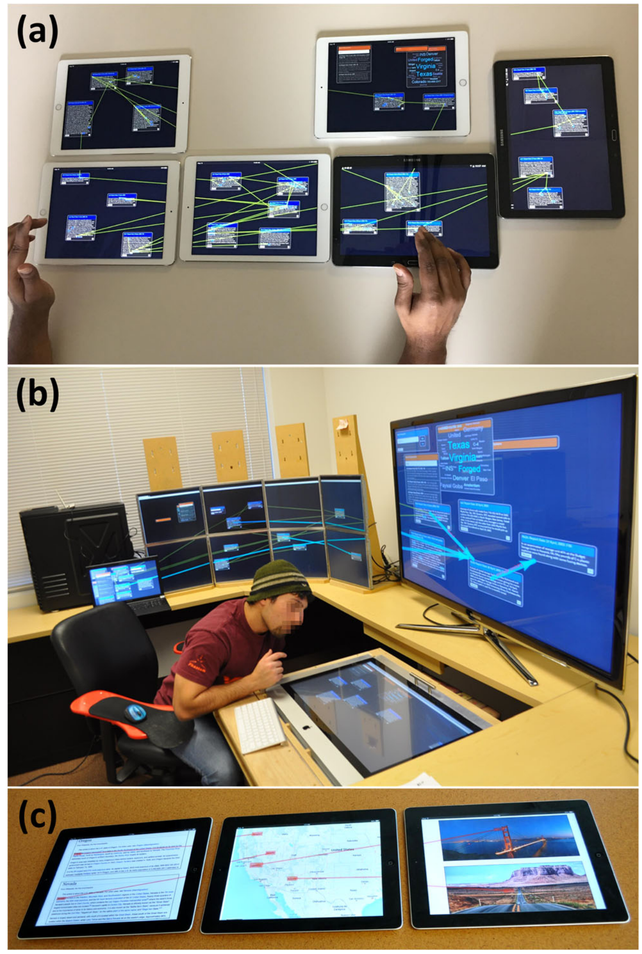

Fig. 1 Spatially Aware Visual Links (SAViL) for sensemaking of multimedia data in various display ecologies. SAViL is designed to visually connect various multimedia data across different displays and devices. **a** A user performs sensemaking of text documents using a number of mobile devices. **b** Screen shot from our evaluation showing how users create semantic layers using SAViL and multiple heterogeneous displays when performing sensemaking. **c** Connecting different types of web documents on different displays through SAViL's DOM-based approach

[20], thereby enabling users to employ the screen spaces of separate displays to create *semantic layers* of documents on multiple displays. By combining readily accessible displays (and thus, the data available on each) through SAViL, it is possible for a display ecology to offer the same space to think benefits as expensive, large, high-resolution displays.

This work contributes to the literature by describing a technique for display ecologies that is expected to increase our understanding of the value of physical space for sensemaking. Specifically, we expect to contribute to the field in the following ways:

> *Cross-display visual representation and interfaces for creating semantic layers*. The primary contribution of this work is to describe a novel cross-display visualization technique and associated tools to support sensemaking

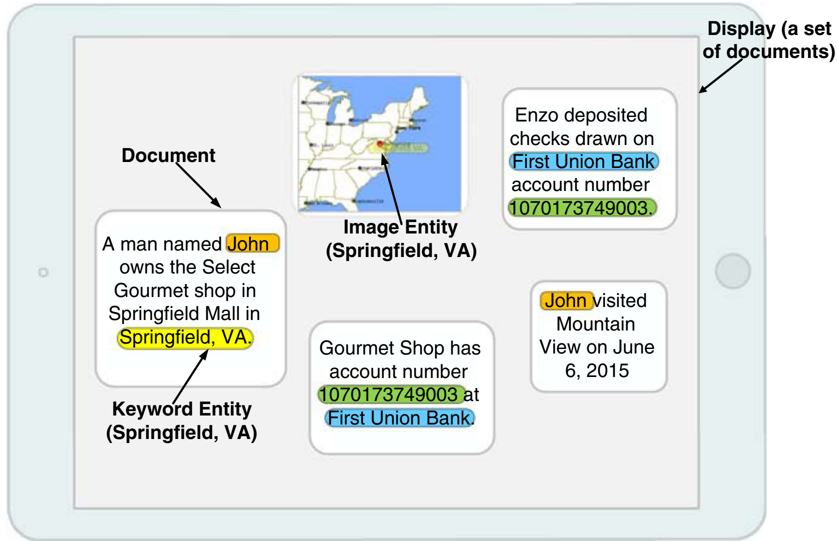

Fig. 2 SAViL's document elements and visual objects on a display. Our target sensemaking workspace on each screen consists of several different visual objects; rounded rectangles represent documents, and color-coded keywords indicate different types of entities in a display. Also, multiple documents form a cluster based on their positions within a display

among heterogeneous multiple displays. The cross-display visual link is designed to connect, direct, and organize document elements of differing levels, located on separate displays and devices, to support sensemaking of text data.

Implementation and system architecture for display ecologies. Our research also contributes to understanding a web architecture that facilitates connecting, visualizing, and synchronizing visual links and data across different displays. Based on this architecture, we present a prototype sensemaking tool. We also present a Document Object Model (DOM)-based approach [21] for linking documents and other multimedia elements visually and spatially across separate independent devices/displays.

Impact on a sensemaking task for creating semantic layers. For the third contribution, we extend prior investigations [20] of sensemaking on large displays to the mixed-display environment, wherein users can utilize large screen real estate from heterogeneous displays for sensemaking. We conducted a qualitative user study to explore the effect of cross-display visual linking on the sensemaking process, which focused on the diverse strategies and processes of creating semantic layers that add meaning based on the position of document elements and displays.

2 Display ecologies and sensemaking

When a user employs heterogeneous displays in a co-functioning ensemble to achieve a goal, they form what is known as a display ecology. The term "display ecology" is defined as an environment of interconnected displays that interact with, and relate to, one another, to assist people in achieving a specific task. We incorporate the "ecology metaphor" as a design concept of how multiple displays mutually interact, support, and collaborate with one another to solve user analysis challenges—rather than designing individual visual and interaction techniques for each single display.

Display ecologies can 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 [7, 22]. In particular, we focus on how display ecologies can better assist users in sensemaking models [2, 3] that improve how people forage for, collect, organize, and produce new knowledge from a document dataset through the use of larger and discretized display space afforded by display ecologies. To this end, we extend the concept of space to think in display ecologies [20, 23]. Specifically, space to think can be facilitated by enhancing the ability to construct semantic layers across different displays, which refer to a cluster of relevant documents based on regions, different topics, timelines, and events or in terms of their importance to the analysis (e.g., Fig. 3). By utilizing semantic layers, users can encode different meanings according to spatial relationships (e.g., proximity, ordering, and alignment of documents) among documents across multiple displays. Semantic layers better enable the user to leverage the larger display space for rapid externalization of cognitive processes during sensemaking. Andrews et al. [20] confirmed that a large, high-resolution display enabled such a phenomenon, in which users utilized the additional screen space to support semantic layers.

SAViL was inspired by prior studies and techniques for space to think in that it focuses on helping users spatially organize information and evidence in order to generate a cohesive hypothesis from multiple displays. Hamilton and Wigdor [4] reported how users employed a number of small mobile displays to conduct a sensemaking task. Specifically, they observed

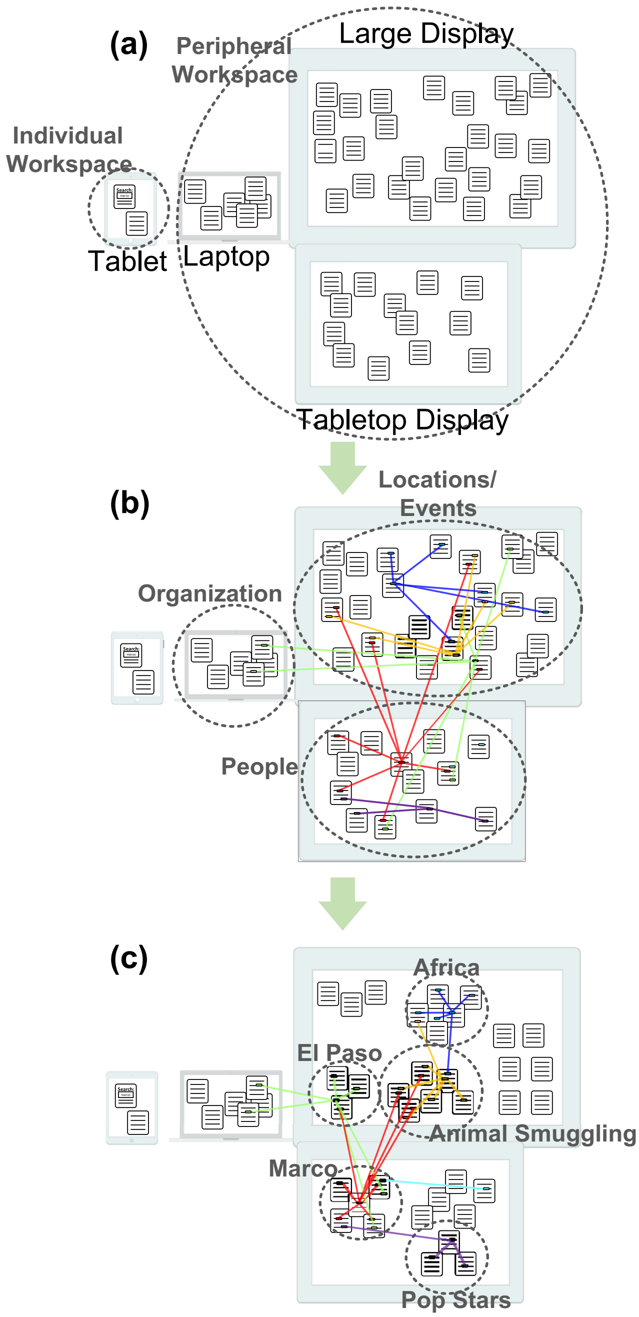

Fig. 3 Example workflow for the semantic layers guided by cross-display visual links. **a** Decide to create a semantic layer of documents by dividing displays into two work zones. **b** Click entities and create an initial layout of visual links that connect the entities across displays. **c** Reorganize clusters and gradually transform initial clusters on displays into more refined ones through using visual links

that users organized the physical devices on a table in order to spatially structure information displayed on those displays. In our prior work [6], we also found that space to think can be extended to mixed-display environments that contain multiple displays of various sizes and form factors, including large displays, desktops, and mobile displays. However, we dentified one critical component (which, to date, remains underexplored) for sensemaking with multiple displays and devices: the ability to identify connections/links between information on different displays. The principal design difference between SAViL and these prior space to think studies are that SAViL facilitates sensemaking tasks via the formation of semantic layers through the aid of visual links among entities and documents located at different displays. In Fig. 2, orange-highlighted entities indicate the same person entity "John," while yellow indicates the location entity "Springfield." These document elements become visual objects shown on each screen. The cross-display visual links visualize the *co-occurrence* of these entities, which shows how the same entities appear in at least two different documents. The use of co-occurring links of the same entities to improve sensemaking has been explored in several sensemaking and visualization research works [18, 19].

In addition, two sensemaking studies based on multiple shared displays are highly relevant to our work. Although they do not specifically focus on forming space to think with multiple displays, both studies investigated how the configuration and type of multiple displays impact the performance and dynamics of teams collaborating on sensemaking tasks. Plaue and Stasko [24] conducted a user study evaluating three spatial display configurations for collaborative sensemaking tasks: single display, side-by-side dual displays, and opposing dual shared displays. Similarly, Wallace et al. [25] investigated the use of three display types for a collaborative sensemaking task: a tabletop display, personal tablets, and a combination of the two. Interestingly, they described the use of "tableaux," which is closely related to SAViL's semantic layers in that a tableaux embodies and externalizes a group's working hypotheses and understanding by spatially forming grids of task slides on one or more displays.

3 Design considerations

To better guide our design of SAViL, we developed three significant design considerations to enable sensemaking tasks with multiple displays. The following design considerations (C1–C3) were guided by existing design considerations for display ecologies [1, 6, 11], as well as by findings from prior related research projects in visual analytics, sensemaking [3], large high-resolution displays [26], and multiple display environments [5].

C1. *Enable space to think with multiple displays*. The sensemaking process can be externalized [6] by distributing and organizing data onto multiple displays. The primary design goal of SAViL is to facilitate spatially organizing and structuring information for the sensemaking process by visually connecting those different displays, enabling users to represent and understand relationships among organized document elements with greater ease.

C2. *Promote utilization of physical space afforded by multiple displays*. A display ecology often implies the composition of various devices in different physical locations. The 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 physical display space provides a better opportunity for physical navigation, which more effectively exploits human spatial senses and embodied cognition [26, 27]. In this respect, SAViL is designed to exploit our physical workspace, where separate displays are located at different places utilizing our physical space.

C3. *Support heterogeneous display ecology.* 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 users have been called upon increasingly to engage with various devices and displays at opportunistic moments. As such, a display ecology may need to be formed with available heterogeneous displays at a moment's notice. Such display ecology scenarios emphasize the smooth reorganization and integration of available displays as an integrated analysis workspace [28, 29]. Thus, SAViL is designed to enable a user to easily distribute visual objects and sensemaking tasks to heterogeneous displays and then connect information according to different levels of details (Fig. 2) across various types of available displays through visual link representations. However, building visual analysis tools based on multiple heterogeneous devices is difficult due to system-imposed constraints. Most notable of these are the heterogeneity of communication protocols and different software and hardware platforms; SAViL supports a web architecture that facilitates managing and synchronizing visual objects and links among multiple displays and devices (Section 7.1), and we also used ad hoc display ecologies in an office for our evaluation (Section 8).

4 Usage scenario

In this section, we describe a sensemaking scenario that illustrates how SAViL can be used for the sensemaking of documents utilizing multiple displays. It also provides a sensemaking workflow, which is based largely on actual results obtained by the authors utilizing SAViL on a specific analytical task that required users to analyze the VAST Challenge 2007 dataset [30].

Noah is a government employee who investigates the illegal possession of endangered animals. He collects data related to the suspicious smuggling of animals, which include 1700 files, encompassing intelligence reports, news articles, charts, and pictures. He initiates the analysis of the collected documents by searching keywords of endangered species on his laptop, which then enables him to locate and open several relevant documents. However, it is difficult for him to synthesize the large volume of information he collects from these diverse data sources.

While analyzing the document dataset with a laptop, he realizes that there are several available displays in his office; thus, Noah decides to utilize his display ecology consisting of one large display, a tabletop display, a tablet, and a laptop. He first decides to create a semantic layer of documents by dividing displays into two work zones: (1) an individual workspace whereby he can read and search documents (tablet) and (2) peripheral zones where relevant documents are clustered. He uses his tablet as the individual workspace and the other displays as peripheral zones (Fig. 3a). He reads documents on his tablet and then starts distributing and organizing documents onto his three available displays, based initially on (1) *people* (tabletop display), (2) *locations* and *events* (large display), and (3) *organizations* (laptop) (Fig. 3b). By simply clicking entities on the documents, Noah creates a layout of visual links that connect entities on different documents located at different displays. He can see the co-occurrence of multiple entities across organized documents at different displays (i.e., the same entities that appear in multiple displays).

Through the aid of visual links among these entities and documents, Noah then forms clusters of documents based on connecting entities or topics (e.g., persons of interest and location of suspected crimes) across displays; he is also able to reorganize document clusters by dragging documents across displays. By utilizing SAViL, he is able to transform his initial clusters on displays into more formal and refined semantic layers (Fig. 3c).

As the size of clusters on each display is increased, Noah seeks to better understand the relationships across different clusters of documents on different displays. On his tabletop, he can first determine how people might be related to each other and to his search of interest. He observes, for example, that one person entity, *Marco*, keeps recurring in many of the documents about a particular animal dealer, as shown on his laptop (where he collected documents related to the organization such as companies and schools). This person of interest, Marco, happens to be a famous pop star who openly espouses conservation wildlife issues.

Noah can be guided to documents related to Marco across multiple displays through the use of visual links (Fig. 3c). He simply follows connected links between the large display (Marco's property ownership) and tabletop (locations/events) to explore other relevant documents. For instance, by reading the linked documents across these displays, Noah is able to identify several co-occurrences of entities in documents and topics related to endangered animals and an animal sanctuary located in El Paso, Texas. Additionally, he can pinpoint other location and event entities that denote activities of interest using the annotation links (e.g., events, importation locations, and even the criminal activities of people associated with Marco).

The linked documents among different document clusters in Noah's display ecology suggest that Marco is actually the behind-the-scenes owner of a suspicious large-animal trafficking operation. Indeed, the seemingly contradictory information makes him more suspicious about Marco, a purported champion of wildlife protection. Because of several visual links to documents related to Marco's association with the exotic animal facility at El Paso (and its ties to the illegal importation of animals from Africa), Noah grows increasingly suspicious of Marco and decides to further investigate whether Marco is smuggling and reselling endangered animals.

5 SAViL overview

The design goal of SAViL is to construct an "integrated visual workspace" over separate displays through visual links. The cross-display visual links are drawn over displays to connect and show relationships between visual objects located on different displays (Fig. 4). In this section, we provide a more detailed description of SAViL's design.

5.1 Connecting visual objects across displays

SAViL is represented with straight lines between visual objects across more than one displays (Fig. 4). Because SAViL enables a user to seamlessly connect visual objects across different displays (e.g., from a laptop to a tablet), it can give the illusion of one continuous workspace over multiple displays, while still maintaining separate workspaces on each display. These cross-display visual links are based on the "partially out of the frame" approach supported by visual techniques for single mobile devices [31, 32]. The theoretical foundation for our cross-display visual links is based on the concept 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 [33].

In each display, the visual objects represent document elements (Fig. 2). Specifically, a document displayed on each screen can be divided into separate visual objects: (1) an *entity* (including a keyword and image region) on a document, (2) a *document*, and (3) a *display* that contains all of the entities and a collection of documents in a screen view. These cross-display links are triggered or created when users click on a specific keyword (or an entity) on a document; a link between these visual objects simply indicates (a) the same entity and (b) different documents sharing the same entities. For instance, if an entity is selected on a document (i.e., a user clicks on it), the selected entity becomes the link source, which means that the links are drawn for all target entities on multiple documents across displays.

Based on these visual objects, SAViL's cross-display visual links can be classified into three types of cross-display visual link approaches: entity link, document link, and display link.

Entity link The entity links are created when users click on a keyword or image region (e.g., a specific area on a map or image). A link between entities is determined by the co-occurrence of the same entity among different documents located on different displays. This approach is designed to help the user develop greater awareness of multiple entities of interest,

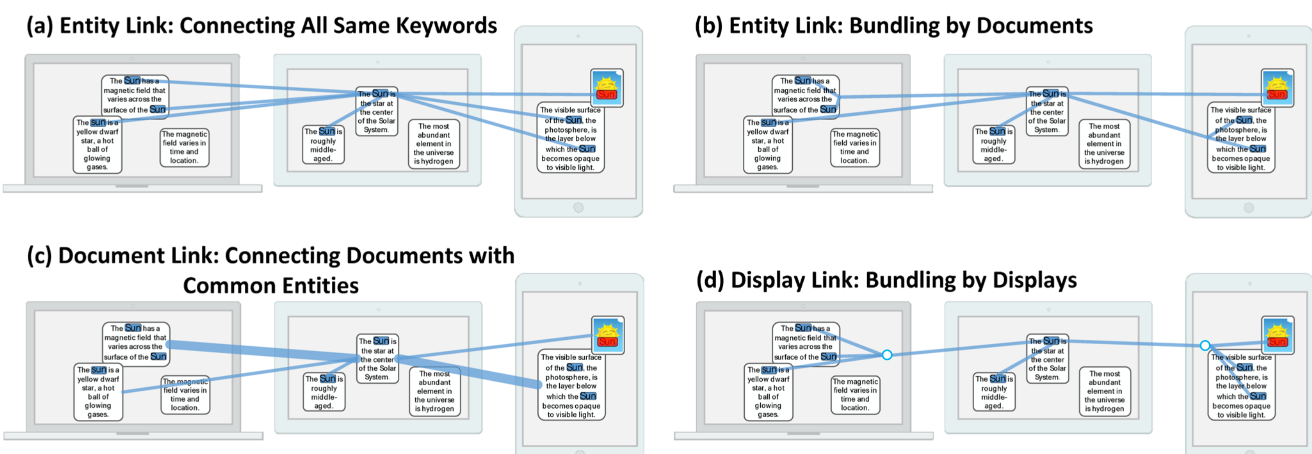

Fig. 4 SAViL cross-display links. Each rounded blue box represents a named entity or keyword. SAViL represents the blue line between documents and between entities across displays. **a** Entity link (connecting all the same keywords). **b** Entity link (bundling by documents). **c** Document link (connecting documents with common entities). **d** Display link (bundling by displays)

which occur in one or more documents (and images) across displays. For example, if a user clicks the entity "sun" on a display, all documents that include that entity on displays are connected directly through the visual links (Fig. 4a). The entity generally indicates a keyword on a document, but it is possible for the user to link between a keyword and its corresponding image region (e.g., a map image in Fig. 2 and a sun image in Fig. 4). However, for an image document, we need to specify an entity for a specific image area manually (see Section 7.2 for details).

However, the entity link can also introduce more visual clutter as the number of target keywords increases. Using a hierarchical relationship between entities and documents, SAViL can bundle entity links associated with the same entities between two documents, instead of connecting all of the same entities separately. SAViL supports simple edge bundling techniques using a hierarchical relationship between entities and documents (Fig. 4b). For this type of link, all of the internal target entities within a document are then connected from the bundling point, which is an intersection point between the bounding box of the document and the links from the same entities. While this type of link reduces the number of links across displays, this approach continues to facilitate the identification of entity occurrences among documents and displays.

Document link Instead of connecting among entities, we can focus on linking between documents. In addition, a single link can be shown between two documents across displays. Through the use of an edge thickness approach, this type of link indicates how many entities are shared between the source and target documents. Varying edge thickness is based on the number of co-occurring entities between the two documents (Fig. 4c).

Display link Based on the bundled document links, these cross-display links are automatically drawn to show the occurrence of entity relationships (i.e., the same entities that appear in at least two displays) and documents across displays. The link source is still a single entity, but the link target becomes any display that contains both the entities and their associated documents. This can further reduce the number of lines across displays (Fig. 4d).

5.2 Annotation links

In addition, users can manually add new links to represent and describe more complex relationships and annotations between two documents across multiple displays (Fig. 5). The other cross-display links described in the previous section focus on representing how the same entities appear in at least two shared documents or displays (i.e., the co-occurrence of entities in two documents or displays). However, there would inevitably be instances when complex semantic relationships would not be able to be represented properly by the co-occurrence of the same entities, such as an alias relationship (between two person entities) or a person's phone number (between a person and phone entities).

Accordingly, we designed a new user interface to allow users to annotate visual links across displays. In contrast to the entity links that automatically connect all entities, the annotating links can only be created manually with a cross-display user interface (UI). For example, in our prototype system (Section 7.3), the user first selects the source document and clicks a link button at the bottom of that document (Fig. 5a, bottom); this brings up a connection anchor icon, as seen in Fig. 5a. To create relationships between two documents on different displays, the user simply drags the anchor across displays (Fig. 5a) and places it on one or more target documents (Fig. 5b). On the anchor, the user can describe a relationship as a link label. When the "Connect Link" button is clicked on the anchor (Fig. 5b, bottom), the overlapped document becomes a link target, and the visual link and its label for the relationship is shown across displays immediately (a cyan arrow link with a text label as shown in Fig. 5c).

5.3 Supporting spatiality of display ecologies

When users work with mobile displays and devices, they can move them from one place to another to re-arrange the analysis workspace as needed for the task at hand. In such cases, users can either manually set the position of each display or change the display topology in a display ecology with the display layout user interface or use the motion-tracking system (see Section 7.1) which can track each display's physical position information automatically.

If a user moves a display to a different location, all of the visual links connected to visual objects in the display are updated according to the new physical position of the display, as shown in Fig. 6. Additionally, if a user changes the spatial layout of documents, all connected links to the associated documents are also followed and updated on each display.

6 SAViL algorithm

SAViL focuses on a tiled display topology/configuration, whereby separate displays are packed together as closely as possible (i.e., within approximately less than 1 cm to 1 m of one another), to create one continuous visualization space, such as tiled display walls [34] (Fig. 7). This display configuration can create one single continuous visualization space through mixing and matching individual displays. This type of a tiled display topology has been broadly and effectively used in multi-display environments for visual analysis and sensemaking tasks. For instance, Plaue and Stasko [24]

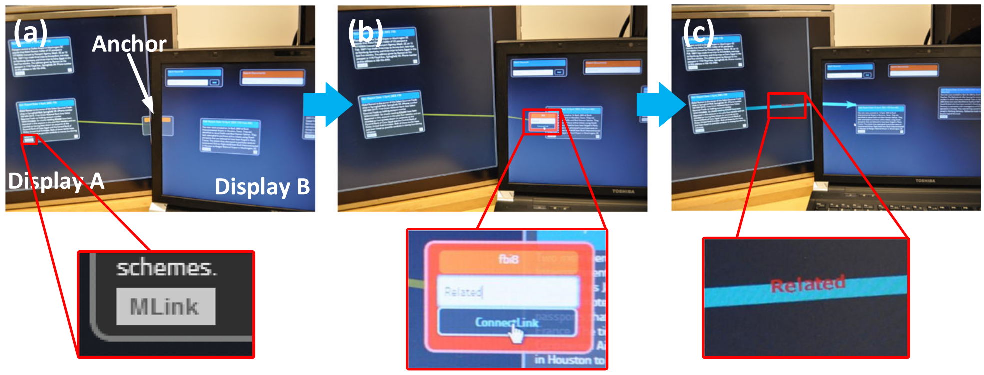

Fig. 5 Annotation links. From left to right: a **a** user drags the anchor across two displays and **b** places it on a target document, and **c** an annotation link labelled "Related" is drawn across the displays

employed a similar display configuration (called the side-by-side configuration) for collaborative sensemaking tasks. They showed that the sensemaking performance in this multi-display configuration allowed collaborating users to identify meaningful information in terms of the total number of insights and inferential links. Additionally, this topology enables simplifying the visual link algorithm, as described in the following section.

6.1 Drawing SAViL over multiple displays

SAViL employs two different coordinate systems to draw and show visual links that connect visual objects among displays (Fig. 7). Specifically, documents and keywords located at different displays are managed in (1) each display's local coordinate and (2) the ecology coordinate that represents a common coordinate shared by all displays in an ecology.

6.1.1 Local coordinate

Each display maintains an independent local coordinate (i.e., the screen coordinate) to render visual objects and visual links within the discrete screen space of each. Additionally, all user interfaces (e.g., mouse or touch interfaces) and their input events (dragging, clicking, or touching) in a display are based on the local coordinates (x, y). However, the distinct positions of the document elements (i.e., visual objects including entities and documents) on the local coordinate of each display cannot be used to draw a common visual link across two different displays (Fig. 7, green regions in each screen). Hence, SAViL needs to support a common coordinate across all different displays.

6.1.2 Ecology coordinate

Ecology coordinate represents the physical space, as well as the common coordinate in which displays are positioned and the cross-display links are drawn globally across different displays (the beige area in Fig. 7). The physical position of each display can first be registered in the ecology coordinate by using the display layout user interface or the motion capture system. Once each display position for the ecology coordinate is determined, the positions of documents and entities in the ecology coordinate can also be calculated relative to each display's position (i.e., each display's top-left corner). Every position of displays and visual objects in the ecology coordinate can be maintained in the central server and shared among different displays (Section 7.1).

In order to render visual links across multiple displays, x, y positions of visual objects in the local coordinate on different displays need to be converted to the ecology coordinate, and a visual link can be then drawn between the target and source

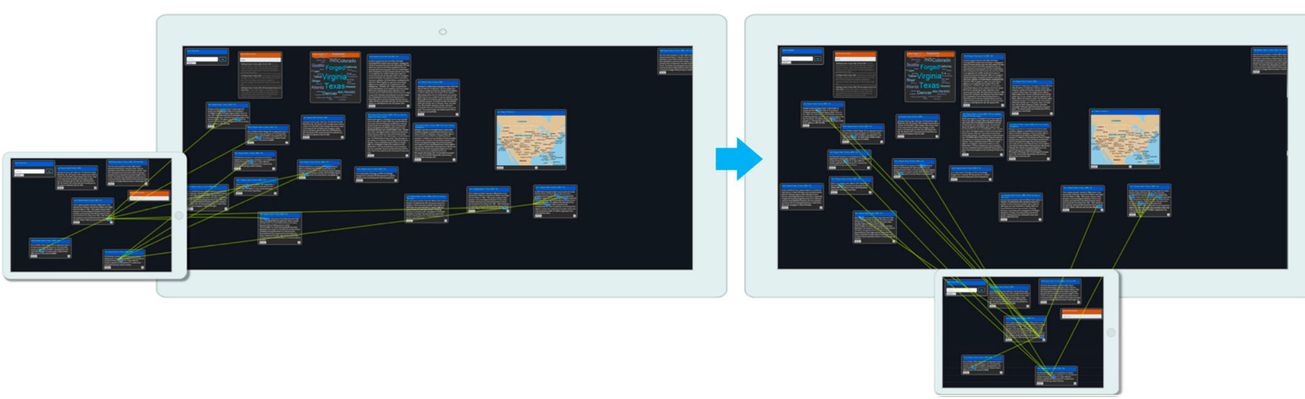

Fig. 6 Support spatially aware links. A small display around a tabletop display is moved to a different location, and the cross-display links keep following the new location of the display

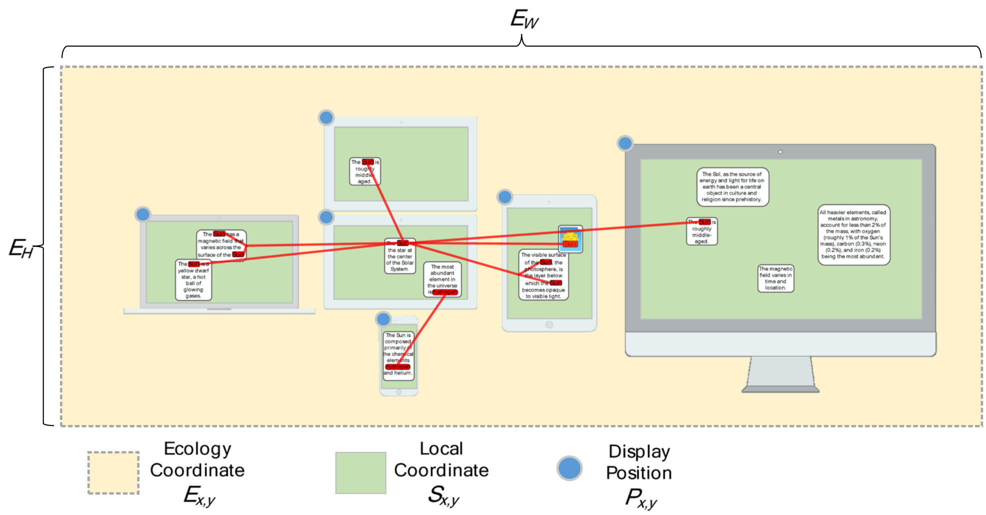

Fig. 7 SAViL coordinates. Visual objects across two displays are managed in two different coordinate systems: the local and ecology coordinates. Each local coordinate is used to render all visual objects on each display, and the ecology coordinate is used to describe the common coordinate among the local coordinates on each screen. The red links represent the visual links drawn on the ecology coordinate.

visual objects located at different displays. To this end, the ecology coordinate of the visual link/object should be converted to the local coordinate to render it. Detailed procedural steps for creating visual links are as follows:

S1. If a user selects a visual object (an entity or document) in a display, a cross-display visual link for the element is virtually drawn between these visual objects in the ecology coordinate. The ecology coordinates (E_x and E_y) of the links, as well as the entities and documents, are stored in and managed in memory on the server through the Artifact manager (Section 7.1.2).

S2. On each associated display in which the visual link needs to be drawn, positions of the cross-display visual links must be converted from the ecology coordinate (the beige region in Fig. 7) to the local screen coordinate of each display (the green regions in displays in Fig. 7).

Given the size (width and height) of the ecology coordinate and the position of a display, a point on the ecology coordinate E_x and E_y can be converted to a point on the local coordinate S_x and S_y and vice versa. Suppose that P_x and P_y are the positions of the current display in the ecology coordinate and that E_W and E_H are the width and height of the ecology coordinate, respectively. The conversion equations from S_x and S_y to a point in the ecology coordinate E_x and E_y is

$$E_x = \left(P_x - \frac{E_\mathrm{W}}{2}\right) + S_x$$

$$E_y = \left(P_y - \frac{E_\mathrm{H}}{2}\right) + S_y$$

S3. Using the local coordinate of each display, the cross-display visual links are then rendered on each associated display in which target or source visual objects are maintained.

S4. If the position or state of each visual object is changed by the user's interaction (e.g., dragging and clicking), the cross-display links are redrawn according to S1, S2, and S3.

This approach also allows us to move different visual objects off the edge (e.g., right, left, up, or down) of one display to another display based on the spatial location of each display; importantly, this can be employed to display any 2D visual shape over multiple displays. For example, based on this approach, we can draw a single visual shape (e.g., circle or rectangle) spanning multiple displays.

6.2 Link width adjustment

Once the visual link has been drawn across two different displays utilizing the above steps, visual properties such as a visual link's line width need to be adapted depending on the properties (e.g., pixel density) of the available display. So, regardless of the pixel density of a different display, the size of visual links should remain uniform. However, 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 line thickness on different displays.

To address this issue, we designed an algorithm to maintain the line width uniformly across different displays. We used pixels per inch (PPI) as the universal measurement standard for pixel density for various displays. In order to maintain a consistent visual link width in inches ($\mathrm{Link_i}$) across different displays, we need to calculate the actual visual link width in pixels ($\mathrm{Link_p}$) in each display based on its pixel density.

Suppose, for example, that one display's physical diagonal is D_i in inches and its width and height resolution of the display are S_W and S_H, respectively. We can calculate the diagonal resolution in pixels, D_p:

$$D_p = \sqrt{S_W^2 + S_H^2}$$

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:

$$Link_p = Link_i \times P_p = Link_i \times \frac{D_p}{D_i}$$

7 SAViL system

In this section, we describe the system architecture and the implementation of SAViL. Additionally, we present a prototype sensemaking environment that implements the SAViL algorithm.

7.1 SAViL architecture

To support sensemaking tasks through visual links in a display ecology, flexible interoperability among heterogeneous displays requires several important capabilities, most notably (1) information transfer between displays, (2) spatial co-awareness between displays, (3) showing and synchronizing visual links and objects on multiple displays, and (4) dynamic device/display membership within a display ecology. In response to these requirements, we present a web-based architecture in which multiple managers store, distribute, and synchronize visual links/objects and user events across different computing devices. The SAViL architecture comprises multiple web clients that run on web browsers and one server. Figure 8 depicts the relationships between the components of the architecture.

7.1.1 The web client

Basically, the view of the client application on each display plays the role of a viewport to the world view (in the ecology coordinate), where all visual objects (documents and entities), as well as their associated links in the ecology coordinate, are rendered and the user can interact with them through user interfaces in a common visualization space. The clients are independent web applications that run at each device in parallel, but they generate the state information for visual links, displays, and visual object properties, including the document ID, document positions, link coordinates, display ID, display positions, and data queries, which is then shared and synchronized among other clients through the server. Once each client receives the state information data from the server, it attempts to render only associated visual objects and visual links within its display, relying on the display position in the ecology coordinate.

7.1.2 The server

Communication among the clients is mediated by the server, which broadcasts the information received from each client to all of the clients in a display ecology. The main role of the server is to watch and synchronize changes associated with the evolving state of visual objects and visual links, as well as to display positions to each client. Specifically, based on different types of state information, the server manages and keeps all of the locations/states of visual objects and displays synchronized via three managers to render coherent visual links across displays and devices. The server consists of three distinct managers, a database system, and a motion-tracking system (Fig. 8).

Device membership manager The SAViL architecture allows users to organize a display ecology flexibly with multiple displays and devices. If a user starts a SAViL client application, the device manager reports the addition of new device to the server. The membership manager allows a user to add or remove devices during the runtime of the application (i.e., one device enters or leaves the current display ecology) and enables displays to update the connected visual links subsequently through sending updated device information to the view manager. Thus, when a display is added or removed, the connected links across displays are immediately updated by adding/removing the new display and its contained documents; the existing visual link layout is also updated by such membership changes. This manager also assigns a unique display ID to each display and maintains a list of display memberships in the ecology.

Artifact manager The artifact manager keeps tracks of the positions of visual objects in the ecology coordinate —e.g., moving, removing, creating, and selecting visual objects. As mentioned, visual objects (e.g., a keyword entity or document) have a unique DOM element ID. This unique ID is used to identify and report the changing status of specific items and document data across multiple displays. This manager oversees such visual object information by storing and sending from/to clients (i.e., displays).

View manager The view manager identifies, saves, and manages each display's physical location (decided manually by the display layout user interface or automatically by motion

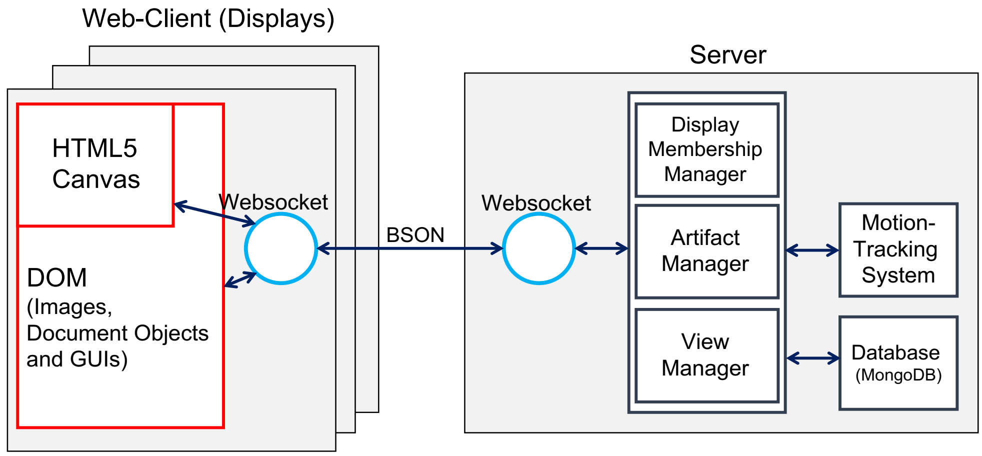

Fig. 8 The SAViL architecture

trackers) in the ecology coordinate and associated views based on each display's resolution. The view manager also allows each client to retrieve the positional information of displays.

Database In addition, the server is connected to the database to store document datasets to process document queries sent from the web clients. For the database server, we employed MongoDB and Mongoose [35] to be able to access the MongoDB commands for searching and reading document datasets and selected entities through the server.

Motion tracking system To automatically identify and track the position of each display in the ecology coordinate, we use six motion capture cameras (OptiTrack Flex 13) and Motive (motion-tracking software and streaming server) [36] with OptiRX [37], which broadcasts motion capture data (i.e., the position of each display) to the SAViL server. In this system, the SAViL server and Motive, which play the role of the streaming server for motion capture data, are maintained on the same machine. To enable tracking of each display, we could attach IR-reflective tape to the bezels of each display; each display has a unique marker arrangement which is then mapped to a rigid body ID. These different rigid body IDs were defined and stored in Motive software and were used to identify and distinguish different displays in the 3D space. Also, we assigned the *pivot* point to a marker at the top-left position, which represents the position of each display (see the blue dots in Fig. 7). However, this motion-tracking feature was not employed for our evaluation, since each display was placed at a fixed position.

7.2 Implementation

One of the design considerations (C3 in Section 3) is how we can link visual objects on heterogeneous displays/computers, thereby allowing users to form a display ecology with the heterogeneous devices available to them. Thus, the SAViL algorithm is implemented as a web application, which has suitable cross-platform capabilities, since the Web supports almost all devices and operating systems (OS) ranging from cell phones to large high-performance computers without any installation process.

As we discussed in Section 7.1, the SAViL architecture comprises multiple web clients and a single server. The functionalities of the web client are implemented in JavaScript, the cross-display visual links are rendered by HTML5 CANVAS and CSS, and the functionalities in the server are implemented with JavaScript and Node.js.

For the client, all of the connected document elements (or visual objects) with SAViL become DOM elements; the client identifies the positions and sizes of these DOM elements on a display in the ecology coordinate. For instance, if a user opens a new document on one display in a display ecology, the location of the document is registered in the ecology coordinate. If the user selects an entity on the document (i.e., by clicking or tapping them on a display), the element is enclosed with a <span> tag and a unique DOM ID is assigned to the element, thereby providing the size and positions of the element automatically. Information on the position and ID of each element is sent as a JSON object to the server; these are later employed to draw links and visual objects across different displays based on the physical positions of each display.

In addition to connecting between entities on text documents, it is possible for the user to link between a keyword entity and its corresponding image region (for instance, the user can connect a keyword *Springfield* to a corresponding region on a map image in Fig. 2). For each image, we specify entities for different image areas manually as the *absolutely positioned element* [38] (a text element with position: absolute in CSS), which is placed at a specific location within an image document. By making the positioned keyword invisible (visibility:hidden in CSS; see an image entity on a map image in Fig. 2) at an image area, we can create a visual link between this image area and a text keyword from documents. However, this requires some additional pre-processing to specify an image region for a keyword entity manually.

On the other hand, in the SAViL architecture, the data between the web clients and server are exchanged in the JSON object through WebSocket, which specifies a document element type (type; "entity" or "document"), display id (displayID), document id (docID), entity name (entityname), position in the local coordinate (x, y), source and target nodes of a visual link (source), etc.

7.3 Prototype sensemaking tool

Based on the SAViL architecture, we created a prototype sensemaking environment that implements SAViL with the goal of demonstrating the effectiveness of SAViL for analysis/sensemaking of text data. This prototype supports all of SAViL's cross-display visual links, as well as provides a suite of additional basic analysis tools to help users explore a large collection of text documents and pictures from the database. The primary interface for the SAViL prototype system is shown in Fig. 9. In this prototype, each display in a display ecology maintains a separate workspace in which different documents or images are searched, opened, and laid out based on related topics. The basic elements of this prototype tool include a word cloud (Fig. 9a) that is able to visualize the frequency of keywords in the dataset; a document search tool (Fig. 9b), which allows users to search documents based on keywords; and the entity highlighting/shoebox interface (Fig. 9c), which allow users to highlight and save important keywords (entities). We used this prototype sensemaking tool to observe spatial organization tasks through SAViL, as described in the evaluation section.

8 Evaluation

In order to evaluate the effectiveness of SAViL for increasing a user's ability to synthesize and explore data from diverse sources, we conducted a qualitative user study.

8.1 Research questions

The main goal of this evaluation was to test whether SAViL would help users create semantic layers and synthesize their hypotheses using a broader spectrum of screens for sensemaking of text data. Specifically, our evaluation focused on seeing how "space to think" [20] in a single large display would extend across multiple displays through the aid of cross-display visual links. This evaluation was guided by the following research questions:

- *Would SAViL help users utilize different types of displays as an integrated sensemaking space?*
- *Would SAViL help users forage for and guide their attention to information on multiple displays?*
- *How can SAViL impact the strategy of creating semantic layers in a display ecology?*

In order to assess evaluation outcomes, we investigated how terrorist plots embedded in a document dataset could be uncovered and represented visually using both SAViL and multiple displays. This evaluation extends prior sensemaking studies that have emphasized the value of space for sensemaking, featuring large high-resolution displays [18, 39], multiple small mobile displays [4], one created from notecards on a tabletop [23], and multiple tiled displays [6]. Particularly, we focused on observing how users engaged with the sensemaking process by forming semantic layers, which enabled them to add semantic meaning (e.g., different

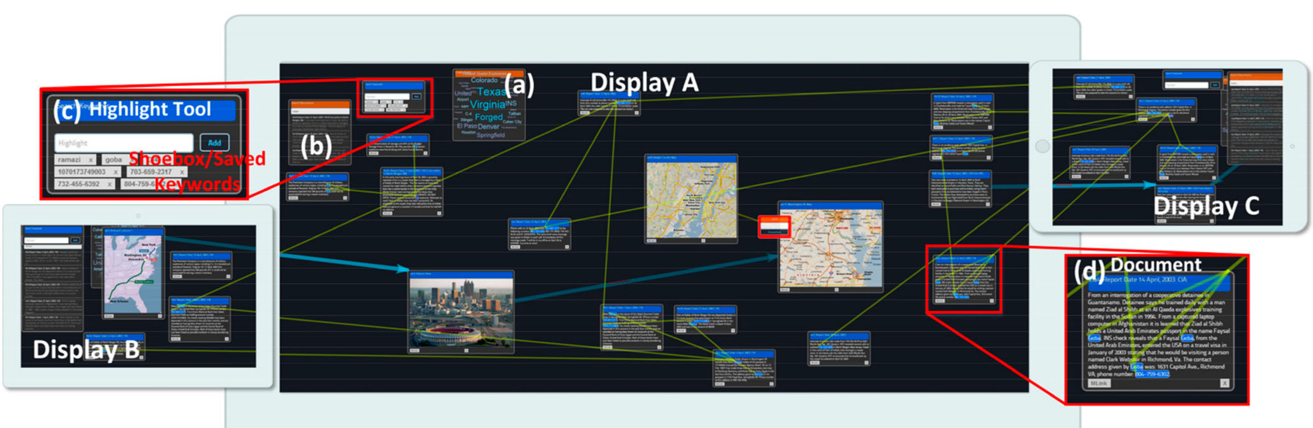

Fig. 9 SAViL with basic document analysis tools. **a** Word cloud. **b** Document search interface. **c** Highlighting and shoeboxing interface to find, highlight, and save entities by entering keywords. **d** Document objects. The light green links indicate the entity links, and arrowed cyan links indicate the annotation links

geographical regions, clusters, piles, or timelines) to the positions or spatial distributions of the displayed visual objects on multiple displays. Thus, we hypothesized that SAViL would encourage participants to form semantic layers of information on multiple displays to understand stories and plots embedded in the dataset, leading to better utilization of displays.

8.2 Evaluation design

The following evaluation methods and the study tasks extended several prior works for evaluating the qualitative aspects of sensemaking based on the concept of space to think [4, 20, 39–41].

Task and dataset In this evaluation, each participant conducted a sensemaking task with our prototype visual analytics with four heterogeneous displays, as shown in Fig. 1b. The sensemaking task was based on the prior work of Andrews et al. [20]—the principal difference is that we employed a display ecology instead of using only a large high-resolution display. The main task for this study asked participants to perform sensemaking tasks through creating semantic layers involving a collection of 41 short (up to 200 words) fictitious textual documents and 22 pictures. These datasets provided evidence of three fictitious terrorist plots and possible associated subplots, in 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. However, we did not score the participant-identified plots but focused on understanding the sensemaking processes underlying the exercise (see Section 9.2.3 for discussion).

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. During a pre-session survey, we confirmed that none of the participants reported familiarity with the use of multiple displays and data analysis. Even though our study included eight students of relatively similar ages as participants, the study did not require special expertise or prior knowledge; it should also be noted that the dataset was originally created for college students [42]. In particular, the study problems to be identified were essentially concerned with typical human actions and motivations. Therefore, we believe we can generalize our results to a broader population cohort and different document datasets.

The participants were randomized and divided into two groups. The only difference between the two groups was whether the visual links were allowed to connect the document elements (visual objects) across different displays.

- The non-cross-display link (NCL) group (U5, U4, U3, U1) was able to use the entity link (Section 5.1) and annotation link (Section 5.2) only *within* each display (i.e., the same entities between two different displays cannot be connected) and all other prototype's analysis tools (word cloud, a document search tool, and the entity highlighting/shoebox interface).
- The cross-display link (CL) group (U8, U7, U6, U2) could use the entity link and annotation link *across* four displays and all other prototype's analysis tools.

Although we focused on analyzing and reporting observations from the CL group, we qualitatively compared the impact of utilizing cross-display visual links on the analytical process and resulting product with a baseline group—the NCL group.

Apparatus For the evaluation, we selected an ad hoc ecology of heterogeneous displays (supporting C3 in Section 3), which is a set of displays that one of the authors actually utilizes in his office. We were interested in illuminating how SAViL can be effective for any type of a heterogeneous display ecology. Participants were provided with a display ecology consisting of four different display types:

1. Fifteen-inch laptop (1366 × 768 resolution) with a keyboard and touchpad connected to a Windows 7 PC;
2. Sixty-inch tiled LCD screen (2 × 4 tiles with total resolution of 5120 × 2160 with a keyboard and mouse on a Windows 7 PC;
3. Forty-inch HDTV (720p; 1280 × 720 resolution) with a keyboard and mouse connected to a Windows 7 PC;
4. Twenty-seven-inch Apple iMac (2560 × 1440, laid horizontally) with a resistive touchscreen on OS X.

For display setups, we set up the spatial layout of the above displays manually with the SAViL prototype's user interface. Figure 1b shows the display layout we used to conduct this evaluation. Based on this spatial layout configuration, each display can be spatially aware of other displays. For example, the visual links should connect the lower edges of the tiled display and HDTV with the top edge of the tabletop.

For user interfaces of each display, the participant was able to share the same mouse and keyboard for all four displays via a mouse/keyboard sharing tool called Synergy [43]. This tool enables the user to move the mouse cursor off the edge (right, left, up, or down) of one display to another display (i.e., different PCs) while sharing the same clipboard. In addition to the single shared mouse and keyboard configured by Synergy, users could employ the touchscreen directly when they interacted with iMac.

Procedures Participants were then given a 10-min tutorial on how to use the prototype system, which focused strictly on system features. After completion of the tutorial, we engaged the participants in the actual experimental session of identifying fictitious terrorist plots and subplots. During the study session, the participants were able to ask any question about the system features. After they completed the 1-h, 15-min, to 2-h session, the user study session concluded with a post-session debriefing to discuss each participant's plots and findings, as well as an individual interview and survey during which the participants were encouraged to use the analytical results on the displays to support their answers. The study results, including the number of documents opened at each display, were measured after the sensemaking session was complete (Table 1).

During the debriefing, each participant was requested to explain their findings and identified plots; we also asked them to discuss regions of clustered documents (i.e., semantic layers) that pertained to specific plots across displays.

Data collection and analysis Throughout the session, video recordings were utilized to capture each study session. The information from the questionnaire and the interviews/debriefing allowed us to identify clusters/semantic layers of text documents and images that had contributed to the formation of any plots. The authors analyzed these collected data based on grounded theory [44] and conducted open coding based on observations and interview notes. By analyzing our multi-sourced data, we sought to better understand the impact of visual linking on the sensemaking tasks utilizing multiple heterogeneous displays. Specifically, the authors discussed each initial code set, conducted coding for the code set together, and determined the final themes related to the processes and strategies of sensemaking through the use of SAViL in the display ecology. These selected themes are reported in the next section.

8.3 Observations and analysis

In this section, we report and discuss our observations during the study sessions, as well as participant feedback for their sensemaking strategies (i.e., creating semantic layers of information) through using both SAViL and multiple displays. All eight participants successfully completed their sensemaking tasks within a 2-h session (NCL mean = 1 h and 20 min and CL mean = 1 h and 31 min), and they created certain clusters of documents across displays to facilitate sensemaking with SAViL. Although each user employed a different analysis approach and different procedures, our observations and interview results indicate that every participant used visual links to identify and understand important documents while creating semantic layers. In the following section, we present a selected subset of their sensemaking strategies and processes through using SAViL. A summary of the study results between the two groups is shown in Table 1.

8.3.1 Enable information foraging and awareness across displays

First, we wanted to learn how SAViL was able to assist users in guiding their attention to relevant documents located on different displays, thereby facilitating cross-display foraging

Table 1 Evaluation results

Group	User	Open documents (actually contributed to synthesizing plots at the end of session)*				No. of screens (actually used to synthesize information)†	No. of distinct plots/subplots identified†	No. of created links (entity links/annotation links)	Elapsed time (h:mm)
		Laptop	Tiled display	TV	Table top				
NCL	U5	6	24	11	0	2	1	18/0	1:19
NCL	U4	0	24	0	0	1	2	36/0	1:28
NCL	U3	0	32	14	0	1	3	25/1	1:18
NCL	U1	0	11	6	0	1	2	4/1	1:15
CL	U8	0	37	8	3	3	3	69/2	1:59
CL	U7	4	20	5	2	2	3	41/0	1:33
CL	U6	9	8	6	8	2	2	15/8	1:11
CL	U2	3	17	10	0	2	2	27/1	1:22

*These results, including the number of documents opened at each display, were measured after the sensemaking session was complete (at the post-session debriefing)

†These results were also based on each participant's explanations during the debriefing/post-session interviews. During the debriefing, each participant was asked to explain their findings and identified plots, and they also explained regions of clustered documents (i.e., semantic structures) that pertained to specific plots on displays

tasks. The cross-display links helped the CL participants maintain awareness of connections between documents on different displays by reminding them visually which documents were linked. The CL participants noted that cross-display links were used to maintain connections between documents scattered on multiple screens, which later aided them in rapidly navigating (going back and forth) among the related documents on different displays. Not surprisingly, the more documents they opened on each display, 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 CL U7 mentioned:

> "After checking other documents on a different display, links make it easier to jump back to the original screen I was working on and refresh my thought process."

CL U8 also made this observation about using visual links in foraging for information flexibly among displays:

> "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…."

After collecting and organizing documents based on different entity types on different displays, the participants connected visual links among the documents (Table 1), including keywords or entities that looked promising, in order to be able to reference those documents more easily later.

Three CL users (U6, U7, and U8) also mentioned that cross-visual links enabled them to locate a set of documents pertaining to multiple topics across different displays because the entity links visually show how documents on different displays were tied together. As an example, U8 (CL) stored documents in one display based on specific *person names*. If a person name had more links with a specific display in which different documents were semantically organized by locations, the user was able to find the specific location-related documents that pertained to that person's name of interest.

In addition, five participants (NCL U1 and U3 and CL U2, U6, and U8) used the annotation links (Table 1) somewhat consistently to reference more complex information between documents, which cannot be guided with links for co-occurrences of entities in multiple documents. If there was information that was related to another document—but the documents did not share the same entities (e.g., an alias or a differently spelled name referring to the same semantic meanings)—they could also be connected with the annotation link. Other participants used annotation linking to tie together pieces of information represented by semantically similar terms (e.g., C4 and explosive). Additionally, the CL participants also annotated documents that contained contradictory information or questions using the annotation link. For example, a CL participant created an annotation link between two documents related to a suspicious person "Webster" and a location "Denver" where an event had occurred, in order to steer himself toward a potential new line of investigation through the use of the annotation "fake name Webster ever visit?"—even though no clear proof of any nefarious activities was apparent on the documents.

8.3.2 Help leveraging multiple displays

Our observations and post-study interviews indicate that the CL participants tended to use more screens for the process of identifying plots through multiple visual links (Table 1). All four participants from the CL group used more displays (mean 3.5/4.0, stdev. 0.57) in comparison to the NCL group (mean 2.0/4.0, stdev. 0.82). Two of the four CL participants used all four displays to organize their documents, and the remaining two participants used three displays. Three of the four CL participants opened at least two text documents on the iMac tabletop; in contrast, none of the NCL group had any interaction with it, except U1. As Table 1 shows, CL participants distributed more documents over more displays. For this result, we wanted to learn why CL participants utilized and added more displays for their sensemaking tasks using SAViL. When questioned during post-session interviews, the CL participants mentioned that the visual links helped them to utilize the display effectively as their spatial organization progressed. Importantly, both the CL and NCL groups cited the two main motivations for using more displays.

First, the cross-display visual links appeared to allow CL participants to maintain awareness of relevant documents across displays, which they added to increase screen space. Three of four CL group members (U2, U7, and U8) added more than one display when they needed to have more screen space for their spatial organizations tasks. They stated that the visual links helped them create new clusters on an additional display space, since the links directly illustrated how new clusters on an additional display are related to existing ones on the other displays.

In particular, all of the participants (both the NCL and CL groups) mentioned that the tabletop display was out of their immediate sight as compared with the other displays (such as the tiled display and HDTV), making them reluctant to use it. This was also the main reason why four NCL users did not have any documents on the tabletop display at the end of their sensemaking sessions. Initially, the NCL participants tried to organize a few (less than three) documents and pictures on the tabletop. However, they did not use any of them to identify the plots. During the debriefing/post-session interviews, the participants were also asked to determine and finalize their document clusters/semantic layers across displays, which involved closing documents on the tabletop that they viewed as “not useful.” Also, three NCL participants stated that it was hard to maintain the context of analysis between the tabletop and other displays, which prompted them to close all documents on the tabletop. For example, although U1 initially collected all picture data on the tabletop, this participant forgot about the pictures and did not use them to identify the plots.

On the other hand, all of the CL participants who used the tabletop stated that they were continuously aware of semantically organized documents on the tabletop through the use of visual links; hence, they could opt to use the tabletop when they needed more display space.

Second, it should also be noted that CL participants (U8 and U2) added an additional display to improve the overall layout of document clusters supported by the visual links across displays. CL U8 stated that this choice was directly motivated by the cross-display links among documents that enabled users to alleviate some of the visual clutter caused by multiple visual links across displays.

> “I found myself having too many links between the TV and tiled displays at the end (specifically horizontal links becoming intermingled); by moving some of documents onto iMac tabletop, I could create more vertical links that were easier to distinguish.”

U2 also mentioned he added a display in his display ecology to improve the overall layout of documents that were connected with visual links. For example, he moved *hub documents* (which are documents with a number of links for entities) to additional displays to improve the overall layout of document clusters and visual links. He mentioned:

> “I started using another display (HDTV) to act as the hub connecting some hypotheses that had previously been on the tiled display.”

However, since NCL participants could not employ the cross-display visual link feature, their spatial organization processes were not related to the overall layout of visual links among the displays. Instead, their motivation for using additional displays mirrored observations associated with a prior multi-display environment study—namely, that adding displays assisted in increasing space for separate entity topics in each screen [6]. Since it was difficult for them to connect and integrate scattered information between two displays, they focused more on using tiled displays to form semantic layers. For instance, one NCL participant (U4) did not utilize any other display except a large tiled display, stating that he did not want to use multiple displays.

8.3.3 Facilitate creating semantic layers over multiple displays

We also wanted to investigate how users enabled the creation of new and meaningful semantic layers or cluster of documents spread over multiple displays through the use of SAViL. The following observations help us understand how SAViL provided participants with space to think, whereby they were able to organize documents into different semantic layers such as regions, clusters, or timelines to reduce the cognitive load in making sense of the content of multiple documents.

As previously mentioned, the spatial organization strategies employed by the NCL group resonate with observations concerning sensemaking and external memory using single large displays [18, 20]. Overall, their spatial layer and organization of documents was confined to the single large display space (mostly on tiled display). For example, the NCL participants performed linking between documents within the tiled display, after which they clustered documents with the links closer together.

On the other hand, although the spatial organization approaches were not consistent for all CL participants, they preferred to use multiple displays for their spatial organization as if they employed a single large display. In particular, through the use of SAViL, they were able to develop hypotheses based on semantic layers across multiple displays. An interesting case was CL U6 (Fig. 10). U6 used the physical positions of the four screens to organize his documents in a semantic way. He organized documents based on chronological order (the orange arrowed line at the bottom of Fig. 10), 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 display) for more recent events. He connected these spatially organized documents via the annotation links (the cyan-colored arrow links in Fig. 10) to describe the chronological order more explicitly. This type of a semantic layer based on multiple displays has not yet been presented in prior studies [4, 6] for sensemaking.

In addition, two other CL group members (U2 and U8) used visual links to guide their semantic layers, as evidenced by the fact that both divided their screens into different spatial categories (topics) across displays (Fig. 11a). For instance, U8 sorted document objects into semantic layers in different displays, such as *people* (suspected perpetrators) on the tiled display, *places* on the HDTV, and *weapons* (or suspicious objects) on the tabletop. He could then connect documents in these

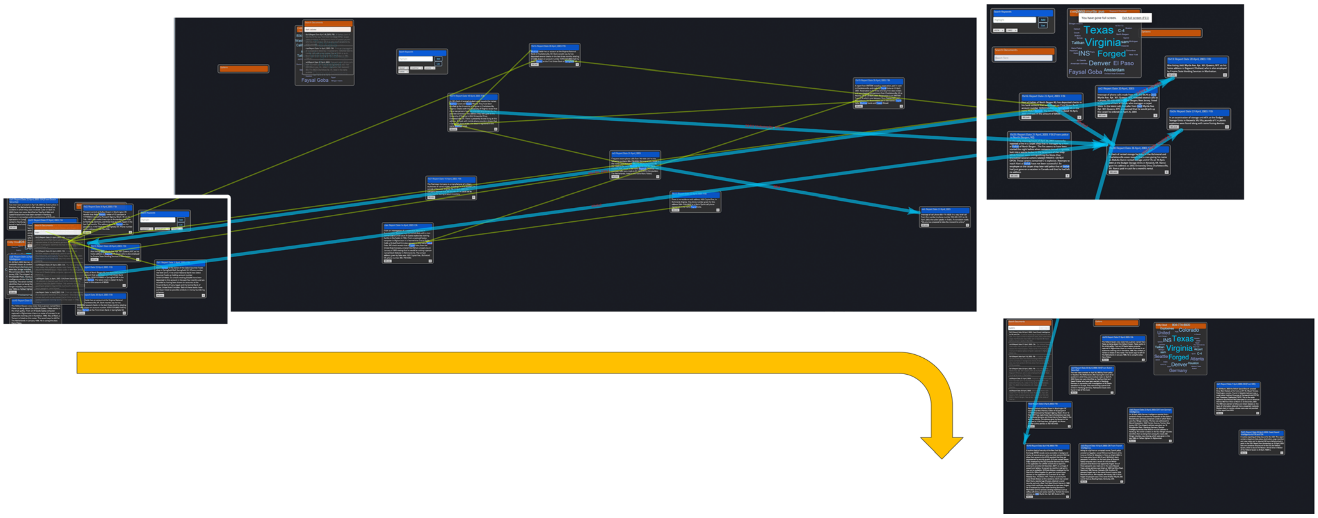

Fig. 10 Exploiting the spatial relationships of displays through the use of semantic layers. Semantic layer U6 organized documents based on chronological order, building a timeline of events with the earliest events on the far-left screen (laptop screen) and then progressing to the rightmost screen (the tabletop device) for more recent events; in so doing, U6 was able to exploit the physical location of displays and manual links. The light green links on each display indicate the entity links, while arrowed cyan links on each display indicate the annotation links. The orange arrowed line at the bottom indicates the chronological order in which U6 spatially organized documents over displays

clusters (i.e., the displays) through the cross-display links (Fig. 11b) to try to gain insights into the relationships and patterns among *people*, *places*, and *weapons*. U2 also sorted documents according to suspicious persons on the tiled displays, while several documents related to location (e.g., NYC, Virginia, and The Netherlands) and suspicious events were clustered on the HDTV.

Interestingly, guided by the cross-display links among the organized documents on different displays, CL U2 moved or opened a few documents between different displays (i.e., different topics) to turn these initial clusters of documents into more formal ones across displays in two different ways:

First, U2 refined or improved existing clusters with the cross-display link between different displays (Fig. 11c). While reading several relevant documents in the accident/event cluster (which were connected by visual links to the same person entity), U2 located additional important documents in connection with the person clusters on the tiled display. Since he determined that some of these documents in the accident cluster were more relevant to his person cluster, he moved and added them from the HDTV to the person cluster on the tiled display.

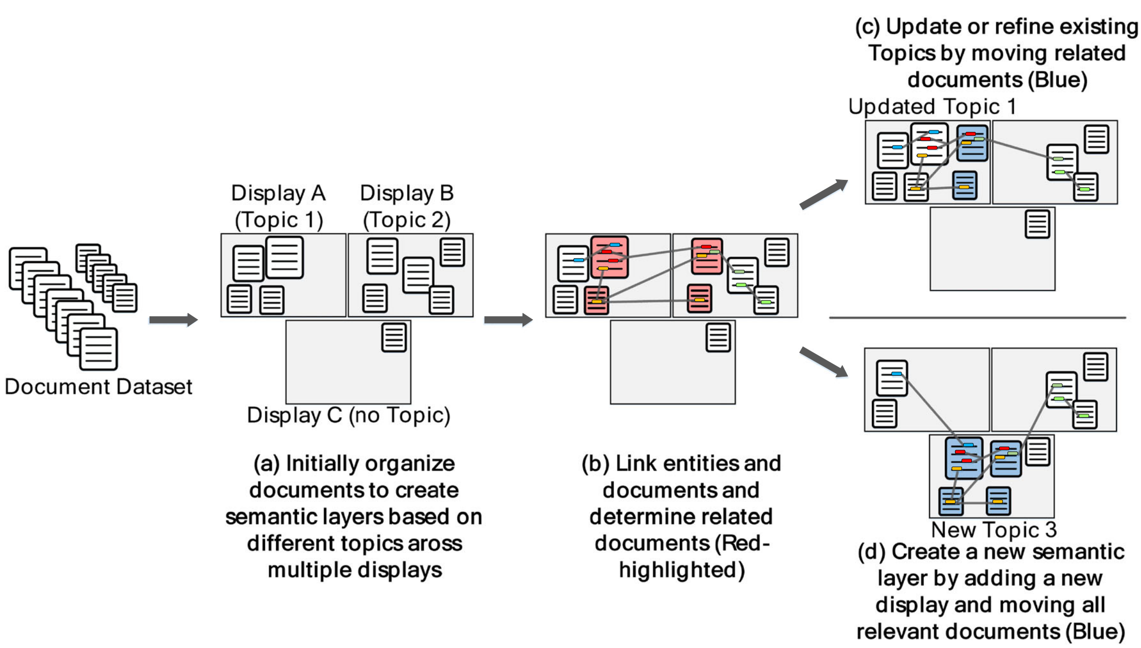

Fig. 11 Two ways to refine and form document clusters guided by SAViL: (a) → (b) → (c) and (a) → (b) → (d)

Additionally, U2 further utilized a cross-display link to create a totally new cluster comprised of information from documents on two displays (Fig. 11d). Specifically, U2 reported that the visual links between two displays facilitated the creation of a new cluster, which was combined of different information related to two topics, since he quickly perceived more visual links among the organized documents across two displays (i.e., locations and persons). Guided by cross-display visual links, he was able to understand how a specific person was related to multiple locations by reading linked documents for two topics; subsequently, he moved or opened three relevant documents on an additional display (laptop) to create a new document cluster. This process led directly to creating a new cluster related to the suspicious person's potential *trip routes* on the laptop screen.

It must be noted, however, that we were not able to observe similar workflow patterns for NCL participants to update semantic layers across more than two displays, since they were not supported by the cross-display visual links.

8.3.4 Promote synthesis of stories across multiple displays

Comparing the CL and NCL participants, it was clear that they synthesized the information pertaining to plots differently. We observed that the main advantage of employing visual links across displays was to enable users to synthesize information from documents on different displays, maintaining the analysis context on each display without switching displays frequently.

All four participants from the CL group and NCL U5 identified at least one plot as a result of synthesizing documents organized on more than one displays. For instance, CL U8 initially organized documents across different displays based on three topics; indeed, one plot that U8 identified was determined by linking documents across displays. Specifically, during the post-session interview, U8 reported that he uncovered the "explosion plan" by linking information on three different screens: five documents from the HDTV, two documents from the tabletop, and two documents from the tiled display—all of which he was able to connect through the use of the entity links. For an analysis workflow, he first organized documents at each display, read the organized documents on each display, and then connected interesting keywords (entities) whenever he found some potentially notable entities by clicking on them in the documents across displays. This process enables him to create a large link visualization spanning three displays. U8 mentioned:

> "I did not move relevant documents from one display to another to create a document cluster because the links explicitly described the relationships of documents placed at different displays."

In contrast, three of the four NCL participants (U1, U3, and U4) reported that they were able to determine each plot by synthesizing the information provided by the documents confined within the tiled screen only—even though they created some document clusters in each display. In other words, when these three NCL participants uncovered one plot, their sensemaking was based solely on organized documents within the tiled display—as opposed to synthesizing documents from multiple displays as the CL group did. For instance, U1 revealed that he identified two different plots related to suspicious terrorist activities, but each plot was synthesized by organized documents only within the tiled display (even though U1 organized six documents on the HDTV). As we expected, this result is potentially attributed to difficulties in semantically connecting document clusters created on separate displays. Specifically, when questioned how they synthesized documents organized on separate displays without utilizing cross-display visual links, NCL participants reported that they used a frequent-switching approach between two different displays. For instance, they first checked the context of plots in documents on one display and then referenced the same entities in other documents on another display—with the goal of understanding and maintaining awareness of relevant documents among different displays without cross-display visual link features.

9 Discussion and future work

In this section, we first revisit the questions that motivated our evaluation and then discuss limitations and propose other interesting avenues for research that manifested themselves over the course of conducting this evaluation.

9.1 Revisiting study questions and results

Our observations during the analysis session, the post-session interviews, and the debriefing session confirm that the cross-display visual links afforded by SAViL do indeed provide novel space to think for sensemaking in a multi-display environment. Based on our evaluation results, we revisit the questions that guided the design of our evaluation.

- Would SAViL help users utilize different types of displays as an integrated sensemaking space?

Throughout our evaluation, we observed that the CL participants were increasingly able to extend their sensemaking space to heterogeneous displays in order to form additional semantic layers. The cross-display visual links, which the participants employed to link among semantically organized documents, helped the participants keep themselves aware of relevant information across physically separate displays—leading to the utilization of more displays (even the tabletop with poor visibility due to its orientation) (Section 8.3.2).

- Would SAViL help users forage for and guide their attention to information on multiple displays?

The CL participants noted that the cross-display links enabled them to maintain awareness and connections between entities and documents scattered on different screens; SAViL also aided them in annotating, locating, and navigating among the related documents from different displays (Section 8.3.1). For instance, while CL participants were conducting their sensemaking tasks, they frequently created visual links between the particular documents organized at each display in order to facilitate referencing those documents easily; they then used the links to locate relevant information on documents organized at different displays as needed. Also, the CL participants annotated specific documents using the annotation links to direct their focus to more complex relationships of documents over different displays, which cannot be easily described through co-occurrences of entities.

- How can SAViL impact the strategy of creating semantic layers in a display ecology?

Guided by the visual links, the participants spatially organized documents across multiple displays or even created unique semantic layers over different displays; they were also able to transform initial random clusters on displays into more formal semantic layers (Section 8.3.3). Also, a single document cluster could be spread over more than two displays through the use of SAViL (Section 8.3.4). Importantly, these semantic layers have not been presented in any prior sensemaking study based on multiple displays. Indeed, we were able to confirm that CL participants formulated plots/subplots as a result of synthesizing information from the clusters they created using multiple displays. In short, the information synthesis conducted by the CL participant was not confined within a single screen. This observation shows that cross-display links could change the ways in which a user conducts a sensemaking task using a display ecology.

9.2 Limitations and future research

We now consider the following future studies to address and understand additional aspects of cross-display visual link and sensemaking in display ecologies.

9.2.1 Many links across displays

The node-link diagram that SAViL employs has been broadly recognized as one of the most intuitive visualization strategies for human observers to understand the relationships between different visual objects [45]. Although our observations confirmed the positive analytical potential in performing sensemaking tasks within a display ecology, a sizable number of cross-visual links among displays resulted in visual clutter and hindered users from viewing clear relationships. In other words, as the number of documents in each display increases, so does the number of cross-display visual links—requiring the user to determine how a specific link connects between entities located on different displays. In support of this supposition, we observed that CL participants spent more time completing their sensemaking task (Table 1). This result may be attributed to having to check a greater number of links and exploring a number of connected entities and documents.

To alleviate this problem, SAViL supports simple edge bundling techniques using a hierarchical relationship based on different levels of detail within documents and displays (Section 5.1). However, our prototype analysis system allowed participants to select only one bundling approach overall, making it difficult for them to use different bundling strategies together. Consequently, all participants focused on using only one type of the visual link (the entity link), which was given as a default link type at the beginning of the study session. This observation highlighted the need to address a new user interface to facilitate using multiple bundling approaches more flexibly. Additionally, we can also employ more advanced bundling approaches to further minimize clutter or the number of edges across different displays. For instance, we take into account several recent edge bundling methods, including force-directed edge bundling [45, 46] and geometry-based edge bundling [47].

9.2.2 Flexible display topologies

SAViL was designed to extend the existing space to think benefits of typical desktop PC environments based on large displays across multiple readily accessible and available displays. Thus, our visual link algorithm focuses more on a tiled display topology [34] which is similar to tiled display walls. In this display configuration, displays tend to be positioned relatively close to each other in a single room; this aggregation of varied screen space increases the overall screen real estate as a single large display, thus enabling users to see more information simultaneously [1]. However, new applications may introduce the likelihood of more variable and flexible display topologies, which would enable two or more mobile devices/displays to be opportunistically combined by physically moving them from one place to another in support of different analysis processes, tasks, and personnel. In particular, such variable display topologies are directly related to appropriately supporting mobile and portable display ecologies. For example, a user may be working with various display topologies or alignment modes that support joint interactions between multiple handheld and wearable displays (smartwatches, smartglasses, etc.) on and around the user [48].

We will continue to determine how cross-display visual links can support mobile display ecologies to enhance both visual analysis and the overall sensemaking process; indeed, supporting such flexible display topologies entails both challenges and opportunities for additional research in supporting visual analysis and user interfaces on multiple displays [49]. For the flexible display topologies, we must consider dynamic view perspective compensation between displays [50], thus enabling users to interact seamlessly with visual links in displays that may continue to change position. As we describe in Section 5.3, such mobile display configurations will also require the use of motion trackers to track and detect the position of multiple displays on an ongoing basis.

9.2.3 Quantitative user studies

In our evaluation, we used a qualitative user study to determine the efficacy of SAViL. Although our current qualitative evaluation method and the number of participants did not focus on examining statistically significant differences between CL and NCL, it has revealed the advantages and effects of utilizing SAViL, in addition to important aspects of the sensemaking processes [20].

We carefully replicated the proven sensemaking task and study methods utilizing multiple displays. Thus, our evaluation methods/tasks enabled us to compare results obtained from using a heterogeneous display ecology for sensemaking with those of performing the same task on large displays [20], multiple homogeneous displays [6], and mobile displays [4]. Similar to these prior studies, we did not seek to determine quantitative scoring for identified plots. Instead, we focused on determining the sensemaking process in terms of forming semantic layers, as well as investigating how users interacted with data and multiple displays utilizing SAViL. Particularly, their process of constructing semantic layers enabled us to understand how users can employ a range of multi-display screen spaces as external memory, as well as utilize physical navigation for their sensemaking of text documents. Thus, the evaluation results described herein can support our belief that utilizing multiple display space as enabled through our SAViL concept potentially leads to better sensemaking performance [18, 20, 26].

However, a limitation of this qualitative study method is that it relies more on subjective data collected from interviews and debriefing sessions. To address this limitation, we intend to develop additional quantitative studies that use statistical hypothesis testing. As an extension of the qualitative sensemaking study described herein, we plan to expand our future research to employ quantitative user studies that will target analysis performance. Importantly, we plan to expand the number of research participants (ideally, to more than 40 subjects) in a quantitative user study to evaluate participant analysis accuracy (score) and efficiency (time) for correctly identified plots, focusing more on how sensemaking outcomes could be improved by SAViL. In addition, we will observe and code statistically significant effects of SAViL against those of using multiple displays without SAViL in making sense of documents. For instance, specific performance effects should include the distance that participants move, the number of cross-device object movements among displays, used screen area, and so forth.

We caution, however, that measuring quantifiable improvements is difficult for sensemaking tasks [20], since there is much variance in participant abilities and study design factors such as participant reading skills and opportunistic and longitudinal characteristics of sensemaking [51, 52].

10 Related work

As aforementioned, SAViL is inspired by existing sensemaking studies and techniques in new display environments and focuses on helping users spatially organize evidence and generate a cohesive hypothesis across multiple displays. In this section, we review prior projects and studies pertaining to off-screen visual techniques, visual links for visual analytics, and multi-display environments and then compare them to SAViL.

10.1 Off-screen and multi-display visualization techniques

The off-screen and multi-display techniques for guiding and linking map information in a single mobile display and maintaining work context in multiple displays are closely related to SAViL.

On the one hand, both SAViL and available off-screen techniques such as Halo [31] and Wedge [32] are based on the concept of amodal completion [32, 33] (see Section 5 for details). The goal of these techniques is to help users infer the virtual off-screen location of the object by employing a number of visual shapes on a visible screen. However, these off-screen visual techniques were not designed for multiple devices—nor have they been used to link and locate data and information on different displays to support sensemaking tasks.

In contrast to these off-screen techniques, SAViL was primarily designed to assist the user in connecting visual objects on physically separate displays at different physical locations. In so doing, SAViL can help the user maintain awareness of scattered information across separate displays. Because SAViL is seamlessly drawn across multiple displays, it gives the illusion of one continuous workspace that utilizes different displays.

On the other hand, SAViL is also related to Dostal et al.'s design space for visual focus-aware/gaze-dependent applications [53, 54]. Specifically, the purpose of SAViL's cross-display link techniques are closely related to their design concept in *maintaining and re-establishing context* in multi-display environments. Similar to their work, our target sensemaking tasks in display ecology are designed to switch between displays in which a different set of documents is organized. However, instead of using gaze-dependent techniques, SAViL helps users keep track of important and relevant information among separate displays through simple link representations.

10.2 Linking information for visual analytics

Visual link representations have been used broadly in visual analytics to identify and present relationships between information objects; examples include Analyst's Notebook [55], VizCept [56], Jigsaw [19], visual links across PC applications [16, 17], and Analysts Workspace (AW) [18]. Similar to SAViL, Jigsaw and AW's visual links are based on the co-occurrence of entities on documents. Additionally, Kang et al. conducted an observational user study using Jigsaw and described how such co-occurrence connections between entities were helpful for the sensemaking process in uncovering an embedded threat [41].

To enhance visual analysis through the use of multiple displays, Chung et al. [1] explored the design considerations associated with how to connect information from different data that is maintained within disparate displays. They presented three different types of connections of information and data—namely, *Overview*, *Explicit*, and *Implicit*—based on how the links can be represented over multiple displays. Basically, SAViL can be categorized as a type of *Explicit* connection of information, since it allows all information to be explicitly connected with visual links spanning multiple displays. Importantly, however, SAViL was also designed to support *Implicit* connections by helping users better understand spatial relationships among semantic layers of document elements organized in different displays.

10.3 Multi-display environments for visual analysis

There are several interactive workspaces designed for supporting visual analysis tasks in multi-display environments. In such environments, multiple displays (mostly large stationary displays) can be located at different places in a room, and types of displays are frequently combined to construct multiple coordinated views in a workspace or laboratory setting. For instance, Zoomable Object-oriented Information Landscape (ZOIL) [5] allows users to freely coordinate documents or visualizations around multiple displays. Using a similar concept, Geyer et al. proposed a multi-display system that enables users to organize individual sketches created on between individual tablets and different displays for sharing and discussion [57]. In these systems, documents or visual objects are related and organized within a common zoomable space, while each display plays a role in supporting a different view for this common visual space.

In addition, the following two multi-display frameworks emphasize creating integrated visual space by expanding visualization views and synchronizing user events on multiple displays. Munin [58] 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). PolyChrome [14] is another framework for multi-device visual analysis applications that augments web-based visualizations across multiple displays and manages event synchronization among them. More recently, several multi-display environments have emphasized annexing spatially aware displays for visual analysis. Schreiner et al. [59] and Langner et al. [29] presented a design concept for tangible visualization views annexing multiple spatially aware mobile displays.

Similar to the above-identified multi-display environments, the design of SAViL highlights the importance of integrating visual objects on different displays into an integrated visual analysis environment, whereby users can exploit multiple display spaces for exploring and making sense of data. However, the target application of SAViL focuses particularly on creating meaningful semantic layers through leveraging multiple discretized screen space to support the sensemaking process.

11 Conclusion

In this article, we presented SAViL, a cross-display visual link that enables a user to connect entities and documents in display ecologies to support sensemaking tasks. SAViL visually represents the links between documents/entities across multiple displays to support the creation of semantic layers of documents for sensemaking.

We conducted a qualitative evaluation in order to test the efficacy of the cross-device links feature for the spatial organization tasks with display ecologies. We observed that SAViL helped users explore, connect, organize, and synthesize data efficiently across multiple heterogeneous displays. Results from our evaluation confirmed that SAViL indeed changes the way multiple displays are perceived. Specifically, the SAViL participants tended to utilize more screen space and create unique semantic layers based on the spatiality of multiple displays as a unified visual workspace. Therefore, results from this investigation lead us to believe that visual links can serve as an important component for transforming separate displays into a display ecology, in which multiple heterogeneous displays function in concert to achieve visual analysis as well as sensemaking.

Acknowledgements We thank Kris Cook (PNNL) for her insightful comments.

Funding information This work was partially supported by grants from the U.S. Department of Defense, NSF IIS-1218346, and New Faculty Research Award from University of Alabama in Huntsville.

References

1. Chung H, North C, Joshi S, Chen J (2015) Four considerations for supporting visual analysis in display ecologies. In: Proc. of IEEE VAST'15, pp 33–40
2. Card SK, Mackinlay JD, Shneiderman B (1999) Readings in information visualization: using vision to think. Morgan Kaufmann Publishers Inc., San Francisco
3. Pirolli P, Card S (2005) Sensemaking processes of intelligence analysts and possible leverage points as identified through cognitive task analysis. In: Proc. of international conference on intelligence analysis, pp 6
4. Hamilton P, Wigdor DJ (2014) Conductor: enabling and understanding cross-device interaction. In: Proc. of ACM CHI'14, pp 2773–2782
5. Jetter H-C, Zöllner M, Gerken J, Reiterer H (2012) Design and implementation of post-WIMP distributed user interfaces with ZOIL. Int J Human Comput Interact 28(11):737–747
6. Chung H, North C, Self JZ, Chu S, Quek F (2014) VisPorter: facilitating information sharing for collaborative sensemaking on multiple displays. Pers Ubiquit Comput 18(5):1169–1186
7. Huang EM, Mynatt ED, Trimble JP (2006) Displays in the wild: understanding the dynamics and evolution of a display ecology. In: Proc. of the 4th international conference on pervasive computing, pp 321–336
8. Marquardt N, Diaz-Marino R, Boring S, Greenberg S (2011) The proximity toolkit: prototyping proxemic interactions in ubiquitous computing ecologies. In : Proc. of ACM UIST'11, pp 315–326
9. Chung H, Cho YJ, Self J, North C (2012) Pixel-oriented Treemap for multiple displays. In: Proc. of IEEE VAST'12, pp 289–290
10. Geyer F, Pfeil U, Höchtl A, Budzinski J, Reiterer H (2011) Designing reality-based interfaces for creative group work. In: Proc of the ACM conference on Creativity and cognition 2011, pp 165–174
11. Chung H, Chu S, Quek F, North C (2013) A comparison of two display ecology models for collaborative sensemaking. In Proc. of ACM PerDis'13, pp 37–42
12. Quigley A, Grubert J (2015) Perceptual and social challenges in body proximate display ecosystems. In: Proc of International Conference on Human-Computer Interaction with Mobile Devices and Services 2015, pp 1168–1174
13. Baldonado MQW, Woodruff A, Kuchinsky A (2000) Guidelines for using multiple views in information visualization, In: Proc. of AVI'00, pp 110–119
14. Badam SK, Elmqvist N (2014) PolyChrome: a cross-device framework for collaborative web visualization. In: Proc. of ITS'14, pp 109–118
15. Wigdor D, Jiang H, Forlines C, Borkin M, Shen C (2009) WeSpace: the design development and deployment of a walk-up and share multi-surface visual collaboration system. In: Proc. of ACM CHI'09, pp 1237–1246
16. Steinberger M, Waldner M, Streit M, Lex A, Schmalstieg D (2011) Context-preserving visual links. IEEE TVCG 17(12):2249–2258
17. Waldner M, Puff W, Lex A, Streit M, Schmalstieg D (2010) Visual links across applications. In: Proc of Graphics Interface 2010, pp 129–136
18. Andrews C, North C (2012) Analyst's workspace: an embodied sensemaking environment for large, high-resolution displays. In: Proc. of IEEE VAST'12, pp 123–131
19. Stasko J, Gorg C, Liu Z (2008) Jigsaw: supporting investigative analysis trhough interactive visualization. Inf Vis 7:118–132
20. Andrews C, Endert A, North C (2010) Space to think: large high-resolution displays for sensemaking. In: Proc. of ACM CHI'10, pp 55–64
21. Nicol G, Wood L, Champion M, Byrne S (2001) Document Object Model (DOM) level 3 core specification. W3C Working Draft 13:1–146
22. Isenberg P, Isenberg T, Hesselmann T, Lee B, Von Zadow U, Tang A (2013) Data visualization on interactive surfaces: a research agenda. IEEE Comput Graph Appl 33(2)
23. Robinson AC (2008) Collaborative synthesis of visual analytic results. In: Proc. of IEEE VAST'08, pp 67–74
24. Plaue C, Stasko J (2009) Presence & placement: exploring the benefits of multiple shared displays on an intellective sensemaking task. In: Proc of the ACM international conference on Supporting group work 2009, pp 179–188
25. Wallace JR, Scott SD, MacGregor CG (2013) Collaborative sensemaking on a digital tabletop and personal tablets: prioritization, comparisons, and tableaux. In: Proc. of ACM CHI'13, pp 3345–3354
26. Ball R, North C, Bowman DA (2007) Move to improve: promoting physical navigation to increase user performance with large displays. In: Proc. of ACM CHI'07, pp 191–200
27. Shupp L, Andrews C, Dickey-Kurdziolek M, Yost B, North C (2009) Shaping the display of the future: the effects of display size and curvature on user performance and insights. Hum Comput Interact 24(1–2):230–272
28. Elmqvist N, Irani P (2013) Ubiquitous analytics: interacting with big data anywhere, anytime.IEEEComputer Mag Computer 46(4): 86–89
29. Langner R, Horak T, Dachselt R (2016) Towards combining mobile devices for visual data exploration. Poster presented at IEEE InfoVis'16
30. Grinstein G, Plaisant C, Laskowski S, O'Connell T, Scholtz J, Whiting AM (2007) VAST 2007 Contest-Blue Iguanodon. In: Proc. of IEEE VAST'07, pp 231–232
31. Baudisch P, Rosenholtz R (2003) Halo: a technique for visualizing off-screen objects. In: Proc. of ACM CHI'03, pp 481–488
32. Gustafson S, Baudisch P, Gutwin C, Irani P (2008) Wedge: clutter-free visualization of off-screen locations. In: Proc. of ACM CHI'08, pp 787–796
33. Sekuler AB, Murray RF (2001) 9-amodal completion: a case study in grouping. Adv Psychol 130:265–293
34. Chung H, Andrews C, North C (2014) A survey of software frameworks for cluster-based large high-resolution displays. IEEE TVCG 20(8):1158–1177
35. LearnBoost (2017) Mongoose. http://mongoosejs.com/
36. NaturalPoint (2011) Optitrack. http://www.naturalpoint.com/optitrack/
37. Astanin S (2017) A pure Python library to receive motion capture data from OptiTrack Streaming Engine. https://bitbucket.org/astanin/python-optirx
38. Consortium WWW (2011) Cascading style sheets level 2 revision 1 (CSS 2.1) specification
39. Endert A, Fiaux P, North C (2012) Semantic interaction for visual text analytics. In: Proc. of ACM CHI'12, pp 473–482

40. Isenberg P, Fisher D, Paul SA, Morris MR, Inkpen K, Czerwinski M (2012) Co-located collaborative visual analytics around a tabletop display. IEEE TVCG 18(5):689–702. https://doi.org/10.1109/tvcg.2011.287
41. Kang Y-A, Gorg C, Stasko J (2009) Evaluating visual analytics systems for investigative analysis: deriving design principles from a case study. In: Proc. of IEEE VAST'09, pp 139–146
42. Hughes F, Schum D (2003) Discovery-proof-choice, the art and science of the process of intelligence analysis-preparing for the future of intelligence analysis. Joint Military Intelligence College, Washington
43. Bolton N (2014) Synergy. http://synergy-foss.org/
44. Strauss A, Corbin J (1990) Basics of qualitative research, vol. 15. Sage, Newbury Park
45. Selassie D, Heller B, Heer J (2011) Divided edge bundling for directional network data. IEEE TVCG 17(12):2354–2363
46. Holten D, Van Wijk JJ (2009) Force-directed edge bundling for graph visualization. In: Computer graphics forum, vol 3, pp 983–990
47. Cui W, Zhou H, Qu H, Wong PC, Li X (2008) Geometry-based edge clustering for graph visualization. IEEE Trans Vis Comput Graph 14(6):1277–1284
48. Grubert J, Heinisch M, Quigley A, Schmalstieg D (2015) Multifi: multi fidelity interaction with displays on and around the body. In: Proc. of ACM CHI'15, pp 3933–3942
49. Grubert J, Kranz M, Quigley A (2015) Design and technology challenges for body proximate display ecosystems. In: Proc. of International Conference on Human-Computer Interaction with Mobile Devices and Services 2015, pp 951–954
50. Pirchheim C, Waldner M, Schmalstieg D (2009) Deskotheque: improved spatial awareness in multi-display environments. In: Proc. of IEEE VR'09, pp 123–126
51. North C (2006) Toward measuring visualization insight. IEEE CG&A 26(3):6–9
52. Chang R, Ziemkiewicz C, Green TM, Ribarsky W (2009) Defining insight for visual analytics. IEEE CG&A 29(2):14–17
53. Dostal J, Kristensson PO, Quigley A (2013) Subtle gaze-dependent techniques for visualising display changes in multi-display environments. In: Proc of the international conference on Intelligent user interfaces 2013, pp 137–148
54. Dostal J, Kristensson PO, Quigley A (2014) Visual focus-aware techniques for visualizing display changes. In: Google Patents
55. IBM: IBM i2 Analyst's Notebook. https://www.ibm.com/us-en/marketplace/analysts-notebook
56. Chung H, Seungwon Y, Massjouni N, Andrews C, Kanna R, North C (2010) VizCept: supporting synchronous collaboration for constructing visualizations in intelligence analysis. In: Proc. of IEEE VAST'10, pp 107–114
57. Geyer F, Jetter HC, Pfeil U, Reiterer H (2010) Collaborative sketching with distributed displays and multimodal interfaces. In: Proc. of ACM ITS'10, pp 259–260
58. Elmqvist N (2011) Munin: a peer-to-peer middleware for ubiquitous visualization spaces. In: Proc. of the 1st workshop on distributed user interfaces 2011, pp 17–20
59. Schreiner M, Rädle R, Jetter H-C, Reiterer H (2015) Connichiwa: a framework for cross-device web applications. In: Proc. of ACM CHI'15, pp 2163–2168