

Collaborative Brushing and Linking for Co-located Visual Analytics of Document Collections

Petra Isenberg^{1,2} and Danyel Fisher²

¹University of Calgary, Canada

²Microsoft Research Redmond, USA

Abstract

Many real-world analysis tasks can benefit from the combined efforts of a group of people. Past research has shown that to design visualizations for collaborative visual analytics tasks, we need to support both individual as well as joint analysis activities. We present Cambiera, a tabletop visual analytics tool that supports individual and collaborative information foraging activities in large text document collections. We define collaborative brushing and linking as an awareness mechanism that enables analysts to follow their own hypotheses during collaborative sessions while still remaining aware of the group's activities. With Cambiera, users are able to collaboratively search through documents, maintaining awareness of each others' work and building on each others' findings.

Categories and Subject Descriptors (according to ACM CCS): Information Interfaces and Presentation [H.5.2]: User Interfaces – Graphical user interfaces (GUI)—Computer Graphics [I.3.6]: Methodology and Techniques – Interaction techniques—

1. Introduction

Activities as part of analytic reasoning tasks are often performed in collaboration. The acquisition, analysis, or interpretation of information, sharing and presentation of analysis results, as well as decision making are all activities that groups often perform together [CR03, TC05, HA08].

In collaborative work, not only can groups process a larger quantity of data than any single person, but collaboration may also improve the quality of the analysis [MK05]. Group members can share their expertise about datasets or analysis methods and can offer and negotiate conflicting hypotheses or interpretations about the data [TC05]. These benefits, though, can be rapidly offset by the overhead of collaboration. Group information sharing typically occurs in phases in which individuals drop in and out of close- and loosely-coupled work [TTP*06, ITC08]. In collaborative visual analytics, for example, group members need to be able to work on their own sub-projects, in which tentative hypotheses can be created, followed, and rejected. The analysts' desire for private work is in tension with their desire to capitalize on the group's shared effort. The group may produce pieces of information that could be useful to the analyst; but the ana-

lyst, immersed in his or her work, may not want to be distracted [BMZ*06, Wea07].



Figure 1: Two users collaborate around Cambiera, implemented on a Microsoft Surface.

In this paper, we explore the idea of trading off individual and group work in a context of collaborative visual analytics. We introduce the concept of *collaborative brushing and linking* as the basis of an implicit visual awareness tool across multiple users. Our work concentrates on situations

in which small groups of people come together in face-to-face meetings to make sense of textual documents. Our goal is to support both the individual and the group in collaborative foraging activities including searching for, reading, and extracting information from visual representations. By providing visual awareness of individual activities we hope to encourage group discussion, negotiation, and shared knowledge building. We use an interactive multi-touch table (Figure 1) as the single shared display on which foraging activities are performed. The particular task we discuss in this paper is that of analysts attempting to find information relevant to reconstruct a story from a large set of textual documents.

2. Collaborative Brushing and Linking

The idea of “brushing and linking” [BMMS91] is now seen as a standard technique in many information visualization contexts. “Brushing” means that different parts of the visualization should be keyed to each other, so that a selection in one area is reflected in other areas; “linking” means that changes in state to a datapoint are reflected in all places that the datapoint is shown. In a classic brushing and linking scenario, we might see the same set of data drawn as both a scatterplot and plotted on a map. When one region on the map is highlighted, the points on the scatterplot that correspond are similarly tagged. Similar to this need for integration and linking of individual views in single-user applications, we define *collaborative brushing and linking* as:

An awareness technique, in which the interactions of one collaborator on a visualization are visible to other collaborators viewing the data items in their own visualizations or views of the data.

This concept adds additional information about the social data analysis process to traditional brushing and linking. Collaborative brushing and linking allows users to communicate implicitly, by sharing activities and progress between visualizations. In asynchronous situations, collaborative brushing and linking essentially becomes interaction-based annotation: it is a way of labeling data with interaction information for other users to consider. In this paper, we consider synchronous work, where collaborative brushing and linking is an awareness tool, communicating current work. Past research (see Section 3.1) has found that many collaborations stumble in trying to ensure that they find an appropriate balance of *sharing knowledge* with each other, ensuring that they can establish *common ground*, and *reduce redundant work*. Additional information necessary for maintaining this balance could include, for example, what data has been looked at and by whom, when and for how long information was read, or it could capture additional information, such as importance and reliability indicators, that collaborators wish to add to the data.

3. Related Work

Our notion of collaborative brushing and linking builds on research both from collaborative visual analytics, and tabletop visualization. We refer to Heer and Agrawala [HA08] for general design consideration for collaborative analysis in distributed settings, and Isenberg and Carpendale for co-located [IC07] settings.

3.1. Collaboration Tradeoffs

The need to support both individuals as well as the group has been identified previously in the Computer-Supported Cooperative Work (CSCW) community. Greenberg and Gutwin [GG98] suggest that task-dependent compromises and additional design work are necessary to balance both needs in distributed collaboration. This principle has also been seen in social software systems: for example, collaborative bookmarking systems also need to be good individual bookmarking systems. Within CSCW, this has been recognized as one aspect of the “Grudin Problem” [Gru88]: users must get value out of the additional work they put into collaborative systems. In many collaborative settings, that value comes from awareness tools that let users know what aspects of the problem their colleagues have covered. To do this, the various participants need to share a common ground: a knowledge of how each other are interacting with the data [CR03, Wea07, HA08]. Several *visualizations* for sharing common ground have been proposed. In distributed, synchronous work, this often requires reconciling distinct views: Brennan et al. [BMZ*06], for instance, merge and fuse distinct private node-link graph representations, in order to show information overlap and common ground in nodes and information items collected and looked at by collaborators. In CoMotion [CR03] objects and events can be explicitly shared by placing them in a shared view (e.g., on a map) and implicitly annotated with interaction history information from different collaborators.

3.2. Tabletop Visualization

Using a digital tabletop as our collaborative workspace has a number of implications on the types of information users can gather about their collaborative activities. First, the immediate co-location allows for multiple channels of communication: we need not assume that all forms of communication are electronic, for example. We can allow users to use informal communication channels by talking or deictic references: they can, in other words, build a shared local context through multiple channels, including informal ones that take advantage of the space. More so, we can allow peripheral awareness: a user can simply pay background attention to another’s work by looking at the shared display or by peripherally observing body movements or gestures. This has been seen in many places, and famously documented for subway workers [HL91]; other research has helped reinforce this importance of locality (and the converse costs

of distance) [SBD99, OO00]. However, using these forms of communication may not be enough to support collaborative visual analytics scenarios. Our system is built to support users working individually on their own sub-projects for at least parts of an analysis session. This may include carefully reading and analyzing specific information artifacts and will limit the attention available to keep track of others' actions. Analysts must, thus, decide when to share information explicitly and potentially interrupt others' thought processes. In order to support this type of analysis work we chose to implement features for *implicit* information sharing between users that may be less disruptive for the individual and also serve as presentation aids to show exploration process and history when users decide to work more closely together.

4. Cambiera: A Tool for Co-located Collaborative Information Foraging

Our tool, Cambiera, is designed for information foraging in the domain of visual analytics [TC05]. Cambiera runs on a Microsoft Surface interactive tabletop (Figure 1) which supports multiple, independent, direct-touch inputs on a 1024 × 768 horizontal display. Cambiera currently supports a small group of up to four analysts, working together.

4.1. Data and Tasks

The task we are interested in exploring is that of intelligence analysts attempting to reconstruct a story from a large set of data. This scenario is not an original one; rather, we have adapted it from a series of ongoing challenges within the Visual Analytics (VAST) community (e. g., [GOL*06]). In the standard VAST scenario, intelligence analysts attempt to decode a large set of documents. Buried within the documents is a single story, spanning multiple documents, that must be reconstructed. For example, a set of newspaper articles may all provide indirect evidence of a crime being planned.

Intelligence analysts often use an exploratory, cyclic process of foraging, evidence gathering, and hypothesis generation. Our system allows users to search through a document collection; it visualizes the resulting documents, and allows direct access to document texts in order for the analysts to get detailed information and to form hypotheses. Evidence can be gathered and arranged on the surface to represent information relevant for a specific hypothesis. Our main dataset comes from the VAST 2006 contest [GOL*06], representing some 1200 fictitious newspaper articles. We will describe two fictional analysts working together on an analysis reasoning task, in order to outline features of our tool. The analysts, Ana and Ben are trying to understand an outbreak of BSE, or Mad Cow Disease, in a farming town; they fear this outbreak may be linked to corruption in city hall. Ben is investigating newspaper articles that mention BSE, while Ana focuses on political events.

4.2. General System Description

The design of Cambiera is intended to take advantage of the special properties of multi-touch, co-located surface computing. Although explicit coordination overhead is lower, since users are in immediate proximity, implicit sharing can relay rich information rapidly. Scott et al. [SCI04] have suggested that allowing users to maintain a personal region of the tabletop allows them to negotiate which work is local and which collaborative. However, rigid boundaries may obstruct collaboration: we want to allow users to flexibly move from loosely-coupled to tightly-coupled collaboration (to use the model by Tang et al. [TTP*06]).

Our design is also influenced by SearchTogether, a collaborative web-search tool [MH07]. In SearchTogether, remote collaborators share query histories and result lists, allowing them to efficiently search the web in groups. Similarly, we want to make it possible for users to see each others' query and reading histories. Cambiera is distinctive as a visual support for information foraging. We want to support information foraging with information visualization, allowing collaborators to explore different perspectives or subgoals from their own views of the data that are only connected by visual awareness cues. In practical terms, this means that we want to allow Ana and Ben to work independently. If they find something in common, or find that they are working on the same documents, they can work more closely.

Specific functionality of Cambiera is described in the following sections. We use these four questions about collaborative awareness to guide our discussion:

- Did another search also find my document?
- Has someone else issued my search?
- Has someone considered the same document?
- Has someone read the same document?

4.3. Presenting Search Results

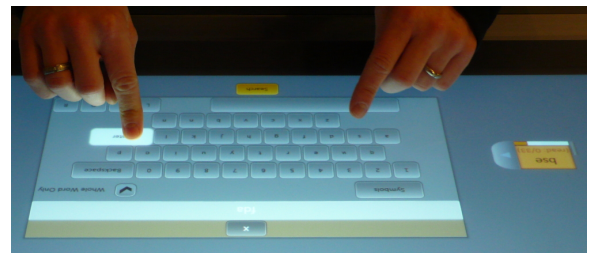


Figure 2: Interaction starts with a search. Each user is assigned a color, which is reflected in the search button (top) and keyboard (bottom).

The core of the system is a search tool for finding documents of interest. In general, a user starts off by pressing a colored search button on their side of the table, which brings up an on-screen keyboard (Figure 2). A search issued from each of these keyboards results in a colored *search*

box, which contains the search results. These search results are first shown minimized (Figure 3, top). Clicking the arrow tab on the right side expands the set to show the individual results returned by the search (Figure 3, bottom). In the expanded view each gray rectangle stands for one of the found documents. All documents are ordered by their publication date. Dark gray bars below each document indicate how many times the search term occurs in the document in order to give the user a hint about which documents may be interesting to explore, reminiscent of ScentedWidgets [WHA07]. On the left side of the widget, we see the search term and a written count of both the search results returned and the number of those that have already been read by any collaborator. A faint bar on the left redundantly encodes the number of documents returned by the search relative to the total number of documents in the collection. The bar is drawn on a log scale to allow easier relative comparison of document count for small result sets. Like all other objects in the user interface, the search boxes can be freely moved around and rotated in the workspace. When a user is done with a search, it can be deleted.




Figure 3: Initial search result overview. One closed search box (top), and one opened search box showing five result details (bottom).

In our scenario, Ben issues a search for FDA (Figure 2); he has already searched for BSE. By looking at the returned *fda* search box, he immediately sees that only five documents were returned and opens the result list to see further details (Figure 3). By looking at the term frequency indicators, he can see that the last two documents include the term *fda* several times and decides to explore those first. Now he wants to find out more about the overlap of his two searches.

4.4. Did another search also find my document?

If a single document was found by two different search terms, then it may be more likely of particular interest. If two investigators are following different paths and stumble on the same results, that source may have particular salience. Our first instance of collaborative brushing and linking enables analysts to follow each others' searches.

In Cambiera, collaborative brushing and linking requires that the system be able to track some information about which user is interacting. While our surface cannot distin-




Figure 4: Color scales to encode search terms. Each analyst's searches receive one hue of their base color.

guish between different people's input (unlike the Diamond-Touch [DL01]), the strategy carried out in this system helps maintain identity. Each user gets a palette of colors that are all variations of one hue (Figure 4); each search gets a distinctive color in that palette. In our example, Ben's searches are orange and always distinctly different from Ana's blue searches. Each document representation that is hit by a search is tagged with the distinct search color. For example, in Figure 3, the background behind the word *bse* received a specific shade of orange to mark this search term. Each document in the search results received a colored stripe in the color of that keyword to indicate that this particular keyword was found in the document. By looking at the stripes (Figure 3), we can see that the first three document in the *fda* set were found both by the results for *bse* and *fda* while the last two were only found by *fda*. This provides a simple way for doing AND/OR searches. The presence of several stripes indicates an AND combination while showing two search representations next to each other allows to see documents that contain either keywords. While our current implementation supports only simple keyword queries, the concept can easily be extended to more sophisticated queries.




Figure 5: Detail-on-demand is shown for the document under the finger. It shows that "bse" also found this document (top-left), a document timestamp, title, and sentences that include the search term (white text, right).

By sliding a finger over the search results, the document under the finger is slightly enlarged for a simple lens effect, and annotated with details. In Figure 5, we see Ben looking at details for timestamp and document title, as well as the text surrounding the one occurrence of the word *fda* in the document text. Left of the title and sentence information Ben can see a summary of search terms that have also found this document, just *bse* in this case. The color is the same as the color of the search representation, and stripes, representing the search term. The size of the word is representative of the number of times this specific search term occurs in the document, reminiscent of TagClouds.

When multiple users synchronously issue searches in the

system the colored stripes in the document representations are updated for all searches. The stripes are persistently linked across all views until a search is removed from the workspace. In our example, Ana has now begun her own searches. She is looking for political connections, trying to understand whether the mayor may have been involved in the BSE outbreak. She wants to see what the mayor and city hall have to say, and so she invokes a search for *luthor*, the name of the mayor, and *city hall*. Ana goes through her city hall list, finding a document that mentions BSE, Luthor, and city hall. The detail-in-context information now includes information for collaborative searches (Figure 6).

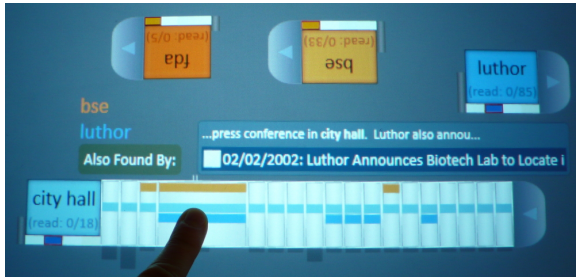


Figure 6: Different base-colored stripes show when searches from other users have found the same documents: Ana has search lists for “city hall” and “luthor”.

4.5. Has someone else issued my search?

As the analysts work, it may be important for them to know that they are currently looking at the same search or at a search a collaborator has previously looked at. Colored rectangles under the search term show which analysts have also previously issued this search (Figure 7). A red border around the box indicates whether there is currently a representation of this search for this user in the workspace. These four base colors used here are not included in the search term scales (Figure 4) in order to make them visually separate. In our example, both Ana and Ben follow their own hypotheses which finally lead them both to issue a search for *mad cow*. The orange and blue rectangles under the search term are both highlighted and receive a red border (Figure 7). The colored stripe on the document itself is also split; this reinforces that the two analysts are both examining this word.

4.6. Has someone considered the same document?

Individual document representations can be dragged up and out of the search results, where they hover in the workspace (Figure 8). The original representation of the document in the search result list is highlighted by a red border to indicate that the document currently resides in the workspace. The basic representations of individual documents contains information on the document title and the publishing date. Stripes are again used to indicate a set of currently colored keywords found in the document (Figure 9, left).




Figure 7: Ana and Ben have both searched for “mad cow.” The search box has both blue and orange marks under it; the stripe that corresponds to the term is split and shows both their colors.



Figure 8: Ana drags a single result up and out of the search box, and so creates a floating representation of a document. Note that this representation shares the striping pattern of the search result.

In our scenario, Ana is interested in a specific document about an event at city hall, she pulls it out of the list (Figure 8) and can immediately see by the red border where the document resides in both the *city hall* and *mad cow* searches. When Ben sees the red indicator for a document in his BSE list, he can reconsider whether he had previously pulled the document in the workspace or start a conversation with Ana about the importance of the document.

4.7. Has someone read the same document?

The minimized representation of a document can be transformed into a full featured document reader by performing a two-handed resize gesture by pulling the document apart (Figure 9, bottom left). In the full document reader, search words are highlighted in their respective color. If the search has been removed from the workspace, past search words are still bolded in black. This allows the user to see how previous searches have touched this document.

When a document is opened in the reader, its counterpart in the search representation receives a slightly darker background to indicate a read access, reminiscent of document read wear [HHWM92]. The darker the background the more times a document has been opened. The color becomes increasingly dark as the document is read more; however, the first color step is distinctive enough that it is immediately visible whether a document has been read. Further reads are

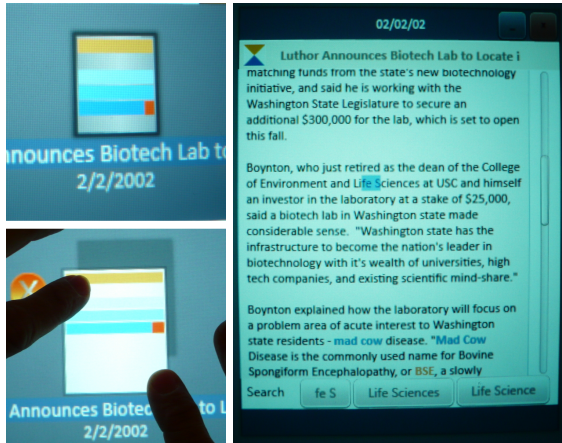


Figure 9: Minimized document representation (top left) and the full document reader (right). The reader is opened by resizing the minimized representation (bottom left).

less finely grained to only reveal whether a document has been read or opened a few or several times as we expect the exact number to be less important. Figure 10 shows an example of two search representations in which individual documents have been read.

In our example, Ana has been concentrated on organizing her documents for a few minutes and was only peripherally aware that Ben had opened a document to read. As the document representations in her set change color, she remains aware of which documents she or Ben have read.




Figure 10: A darker background for individual documents indicate that a document has been opened in the document reader. A darker color indicate repeated document access.

Information about what documents have been read and by whom is important to information analysis [Rob08]. In addition to the gray-scale encoding of read access in the document representations, we also designed a glyph to encode *who* had previously read a document. Figure 11 shows the basic glyph on the left. To its right are examples of the glyph that shows that both the orange and blue user have read the document; their respective triangles are now opaque. The glyph is oriented to match the locations of the search buttons for each user in the workspace; it is rotated appropriately to reflect the viewpoint of the owner of the object the glyph

is embedded in. Figure 11(b) shows the glyph embedded in the detail-on-demand information of the search box, in Figure 11(c) it is embedded in the title of the document reader, and in Figure 11(d) it is embedded next to the title of a document representation floating in the workspace. This last icon is also rotated towards the orange analyst.



(a) Glyph. (b) In detail-on- (c) In document (d) In document demand. reader. title.

Figure 11: Icon representing who read a document. Each triangle stands for one analyst. The icon is embedded in three places. The three examples show documents that have been read by both the blue and orange analyst.

4.8. Sharing Results

During their work, analysts may realize that one particular document or search may be more relevant for one of their collaborators. In this case, the documents and searches can easily be handed over by passing them to the respective colleague who can then drop it on his or her search button. This recolors the item and any interaction with it will now be issued from its new color. In our case, Ben passes documents related to the politics at city hall to Ana while she does the same for Ben when she comes across something relevant to his BSE investigation.

4.9. Searches from Documents

We bring the interface full circle by providing functionality for analysts to issue new searches from the document reader, allowing them to pursue hypotheses through a chain of documents. On a finger's touch, the whole word under the finger is selected; dragging the finger selects a range of letters. In the bottom of the document reader (Figure 9, right) the user is then presented with several search options: both the exact selection and whole words contained in the selection. In the figure, the user has touched the word "Life" and moved his or her finger towards the word "Sciences." The search suggestions for this selection include both the exact match "fe S", the contained words "Life Sciences" or the singular "Life Science."

5. Initial Evaluation and Future Work

In this section we report on initial findings from one security specialist and two pairs of researchers using our system. To elicit initial comments and observations we first demonstrated the tool to a security specialist at our company, encouraging him to work with it for 15 minutes. He found that it was easy to query the dataset, and felt that the direct interaction with the visual results allowed him to easily understand what he had found. He rapidly suggested a number

of contexts in which the tool would be useful, ranging from time-critical search and analysis scenarios to large dataset exploration. He also suggested that the tool could be fruitfully adapted to distributed situations.

To get a sense about how our tool supports the visual analytics scenario discussed throughout the paper, we conducted informal observations with two pairs of researchers using the aforementioned VAST contest dataset. Participants were given a 15 minute training session, worked on the task for 40 minutes, and then discussed their experience with us. Both groups had no difficulty reading or understanding the collaborative brushing and linking features. We saw these features used as intended: for broadening coverage of the information search, as a highlight for convergence, and as indicators to start an exploration. Interestingly, different members took advantage of the awareness features in different ways. One researcher was particularly concerned about common ground, and so made sure to read all documents that his partner had read. Another issued a broad set of queries, and then looked for overlapping terms between them. While the participants reported that they spent time monitoring each others' progress, they worked mostly in silence, only speaking occasionally to confirm a finding or discuss a document. Due to the given timeframe participants did not solve the complete challenge but were able to find partial solutions.

Based on their experience working with the tool, participants asked for an option to bookmark or tag documents as a more explicit way of categorizing their findings. We used a relatively small surface table for this task which caused participants to lose overview of documents in the workspace that may be hidden under an opened document. For the size of table we used we recommend no more than two concurrent users for this task unless reading is performed on an external display. In terms of scalability, our system is currently most optimal for result sets containing up to 50 entries due to the display size. Each document in the dataset can be brushed with up to 24 concurrent searches due to the chosen color scales that hold 6 colors per user. We have addressed this scalability issue by fading out old searches to gray – should more than 6 searches be added per person – and by adding interactive features that can re-color an expired search. Using other visual variables than hue and saturation to distinguish searches and adding ways to refine and combine queries would be another way to address this issue.

We expect that our collaborative brushing and linking features could easily be used on larger surfaces, in multi-display environment (MDEs), or with small modifications in distributed scenarios. We are currently experimenting with a larger surface table and hope to conduct more formal studies in this environment. As a next step we intend to include more direct support for hypothesis formation and evidence gathering. This will include adding features as requested by the participants but also include techniques to summarize or visualize found documents or document content according to time, rank, category, or other semantic relationships.

6. Conclusions

Collaborative brushing and linking is an awareness mechanism that enables users to work independently during collaborative sessions. We note that there are several important differences from interactive brushing and linking. Unlike interactive brushing and linking, the collaborative form links views through social information about items being worked on. Also, interactive brushing and linking can be ephemeral: a user can select a node on one chart, and see the highlight propagate to another instantly. In contrast, in an awareness tool, collaborators may not notice changes instantly; changes may have to be persistent or propagate slowly. Finally, interactive brushing and linking usually connects distinct visualizations with the same underlying data. To help maintain common ground, Cambiera uses the same type of representation throughout. While it may be possible to use different representations, we have not yet experimented with that.

In our design, we have emphasized persistent colorings in order for users to share common ground. In previous sections, we have outlined four different forms of collaborative brushing and linking:

- search stripes, to help users see other search terms,
- document read wear to show what documents have been read and by whom,
- red highlights around documents to cue that the document is visible in the workspace, and
- search boxes which show who has repeated the same search.

Cambiera is designed for a synchronous, co-located surface computing system. While the concepts of shared awareness information are well known for distributed settings, co-located tasks can benefit from unobtrusive awareness cues. Highlighting parts of data that other analysts have seen can help lead them toward documents of particular interest, help validate each others' results, avoid redundant work or, generally, prompt conversations about the data. The visualizations and representations of this shared information are certainly extensible to distributed or asynchronous systems, and would work very similarly: cues might help other users know what documents have been examined, and by whom.

In summary, Cambiera is a system for information foraging activities on a multi-touch tabletop display. Collaborative brushing and linking allows users to maintain common ground and awareness as they work, loosely coupled, on visual analytics task. Figure 12 gives a final overview of the workspace during an analysis session.

7. Acknowledgements

We would like to thank the following researchers for their valuable input: Kori Inkpen, Chris North, Tim Dwyer, Merrie Morris, Hjordve Benko, and Bongshin Lee.

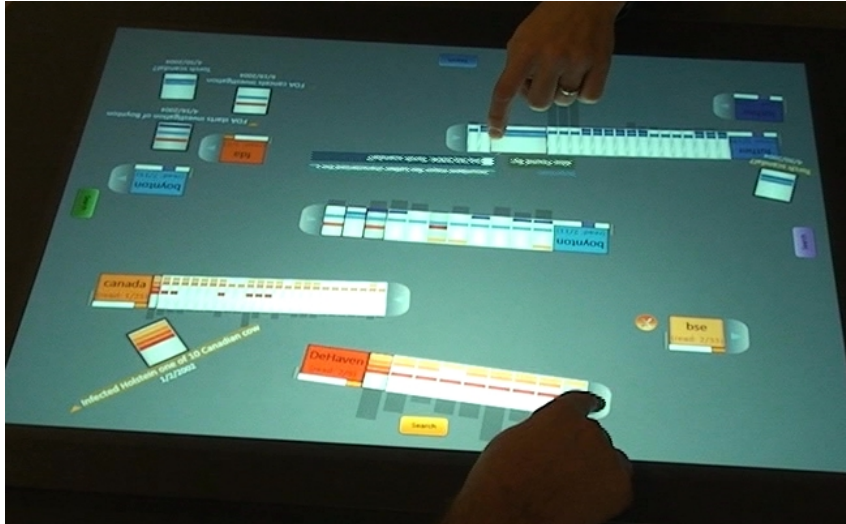


Figure 12: Overview of the workspace during an analysis session. Both analysts have arranged several search boxes and documents in the space related to their current hypotheses.

References

- [BMMS91] BUJA A., McDONALD J. A., MICHALAK J., STUETZLE W.: Interactive Data Visualization using Focusing and Linking. In *Proc. of Vis* (Los Alamitos, CA, USA, 1991), IEEE Comp. Society, pp. 156–163.
- [BMZ*06] BRENNAN S. E., MUELLER K., ZELINSKY G., RAMAKRISHNAN I., WARREN D. S., KAUFMAN A.: Towards a Multi-Analyst, Collaborative Framework for Visual Analytics. In *Proc. of VAST* (Los Alamitos, CA, USA, 2006), IEEE Comp. Society, pp. 129–136.
- [CR03] CHUAH M. C., ROTH S. F.: Visualizing Common Ground. In *Proc. of Information Visualization (IV)* (Los Alamitos, CA, USA, 2003), IEEE Comp. Society, pp. 365–372.
- [DL01] DIETZ P., LEIGHT D.: DiamondTouch: A Multi-User Touch Technology. In *Proc. of UIST* (New York, NY, USA, 2001), ACM Press, pp. 219–226.
- [GG98] GUTWIN C., GREENBERG S.: Design for Individuals, Design for Groups: Tradeoffs between Power and Workspace Awareness. In *Proc. of CSCW* (New York, USA, 1998), ACM Press, pp. 207–216.
- [GOL*06] GRINSTEIN G., O’CONNELL T., LASKOWSKI S., PLAISANT C., SCHOLTZ J., WHITING M.: VAST 2006 Contest—A Tale of Alderwood. In *Proc. of VAST* (Los Alamitos, CA, USA, 2006), IEEE Comp. Society, pp. 215–216.
- [Gru88] GRUDIN J.: Why CSCW Applications Fail: Problems in the Design and Evaluation of Organizational Interfaces. In *Proc. of CSCW* (New York, USA, 1988), ACM Press, pp. 85–93.
- [HA08] HEER J., AGRAWALA M.: Design Considerations for Collaborative Visual Analytics. *Information Visualization* 7, 1 (2008), 49–62.
- [HHWM92] HILL W. C., HOLLAN J. D., WROBLEWSKI D., MCCANDLESS T.: Edit Wear and Read Wear. In *Proc. of CHI* (NY, USA, 1992), ACM Press, pp. 3–9.
- [HL91] HEATH C., LUFF P.: Collaborative Activity and Technological Design: Task Coordination in London Underground Control Rooms. In *Proc. of ECSCW* (Norwell, USA, 1991), Kluwer Academic Publishers, pp. 65–80.
- [IC07] ISENBERG P., CARPENDALE S.: Interactive Tree Comparison for Co-located Collaborative Information Visualization. *IEEE Transactions on Visualization and Computer Graphics* 13, 6 (Nov./Dec. 2007), 1232–1239.
- [ITC08] ISENBERG P., TANG A., CARPENDALE S.: An Exploratory Study of Visual Information Analysis. In *Proc. of CHI* (New York, NY, USA, 2008), ACM Press, pp. 1217–1226.
- [MH07] MORRIS M. R., HORVITZ E.: SearchTogether: An Interface for Collaborative Web Search. In *Proc. of UIST* (New York, NY, USA, 2007), ACM Press, pp. 3–12.
- [MK05] MARK G., KOBASA A.: The Effects of Collaboration and System Transparency on CIVE usage: An Empirical Study and Model. *Presence* 14, 1 (February 2005), 60–80.
- [OO00] OLSON G. M., OLSON J. S.: Distance Matters. *Human-Computer Interaction* 15, 2 & 3 (2000), 139–178.
- [Rob08] ROBINSON A.: Collaborative synthesis of visual analytic results. In *Proc. of VAST* (Los Alamitos, CA, USA, 2008), pp. 67–74.
- [SBD99] STEWART J., BEDERSON B. B., DRUIN A.: Single Display Groupware: A Model for Co-present Collaboration. In *Proc. of CHI* (New York, NY, USA, 1999), ACM Press, pp. 286–293.
- [SCI04] SCOTT S. D., CARPENDALE M. S. T., INKPEN K. M.: Territoriality in Collaborative Tabletop Workspaces. In *Proc. of CSCW* (New York, NY, USA, 2004), ACM Press, pp. 294–303.
- [TC05] THOMAS J. J., COOK K. A. (Eds.): *Illuminating the Path: The Research and Development Agenda for Visual Analytics*. NVAC, August 2005.
- [TTP*06] TANG A., TORY M., PO B., NEUMANN P., CARPENDALE S.: Collaborative Coupling over Tabletop Displays. In *Proc. of CHI* (New York, 2006), ACM Press, pp. 1181–1290.
- [Wea07] WEAVER C.: Is Coordination a Means to Collaboration? Panel Statement, Proc. of Coordinated & Multiple Views in Exploratory Visualization, 2007. Zürich, SZ.
- [WHA07] WILLETT W., HEER J., AGRAWALA M.: Scented Widgets: Improving Navigation Cues with Embedded Visualizations. *IEEE Transactions on Visualization and Computer Graphics* 13, 6 (Sept./Oct. 2007), 1129–1136.