CLIP: A Visual Thinking Space to Support Collaborative Sensemaking and Reasoning *

Narges Mahyar
University of Victoria
Department of Computer Science
nmahyar@uvic.ca

Melanie Tory
University of Victoria
Department of Computer Science
mtory@uvic.ca

ABSTRACT

We explore how linked common work (LCW) can be employed within a 'collaborative thinking space', to support synchronous collaborative sensemaking in Visual Analytics. Collaborative thinking spaces provide an environment for analysts to record, organize, share and connect results. Our tool, CLIP, extends earlier thinking spaces by integrating LCW features that reveal relationships between collaborators' findings. Collaborators' externalizations can be integrated into a single view, with each person's contributions identified using colour coding. Results of a user study demonstrated that LCW in CLIP significantly improved analytic outcomes at a collaborative intelligence task.

Keywords

Sensemaking; visual analytics; externalization; linked common work; collaborative thinking space.

1. INTRODUCTION

When a group of analysts investigate data, they go through an iterative process to generate and evaluate hypotheses, find evidence, and share findings amongst members of the investigative team. This is known as the *sensemaking process*. Collaborative visual analytics tools can facilitate this process by enabling team members to record, organize and share their results [2, 7, 10, 17].

Our research explores the use of *linked common work*(LCW) to facilitate synchronous collaborative sensemaking. With LCW, common work elements such as similar findings are automatically discovered, linked, and visually shared among the group. We built this technique within a 'collaborative thinking space' that enables analysts to record, organize and schematize their externalizations (e.g., notes representing findings and hypotheses). LCW reveals similarities in people's externalizations, enabling analysts to acquire awareness of each other's findings, hypotheses, and evidence. Moreover, each individual analyst can review and merge others' work from within his/her work space. Our results demonstrate that applying LCW to externalizations, and providing the ability to integrate collaborators' findings together within one view, noticeably improve team awareness, coordination, communication, and analytic outcomes.

Recently there has been a growth of tools to support various aspects of the sensemaking process. Sensemaking has been described as involving two iterative loops [14]: the information foraging loop involves searching, reading, filtering, and extracting information, whereas the sensemaking loop involves iteratively developing a mental model, forming and evaluating hypotheses, and publishing the results. Tool support is needed to help analysts with both of these phases. For individual work, many tools have been developed to support both phases of sensemaking (e.g., [1, 8, 18]). In contrast, for collaborative sensemaking much less has been done, and most tools that do exist focus on the information foraging loop [9] or asynchronous collaboration [3, 17]. Most notably, Cambiera [9] introduced the concept of collaborative brushing and linking as a way of maintaining awareness of a collaborator's information foraging activities.

The notion of LCW closely resembles *collaborative brushing and linking* [9] in which actions of one collaborator on a visualization are visible to other collaborators through their own views. However, collaborative brushing and linking was only applied to search queries and retrieved documents and did not cover externalizations. It also focused on supporting only information foraging. In contrast, our work facilitates later stages of the collaborative sensemaking process (i.e., the sensemaking loop), by applying the linking concept to people's externalizations.

Here we describe the design and implementation of CLIP, a collaborative visual thinking space. CLIP allows analysts to record their findings in the form of a node-link graph and timeline, add supporting evidence, and add free form text to record hypotheses, questions, to-do-list, etc. Most importantly, CLIP incorporates LCW to relate and integrate the findings of different collaborators. In this paper we focus on CLIP's design. However, we have also conducted a user study to evaluate CLIP's LCW features; results of the study demonstrated that LCW can lead to more effective group coordination and better analytic outcomes.

2. RELATED WORK

Here we summarize guidelines on how to support collaborative sensemaking, and how those relate to CLIP's design.

2.1 Record-keeping and Schematizing

Mahyar et al. [12] demonstrated the critical importance of record-keeping during collaborative analysis, and the lack of support for this process in current visualization tools. Schematizing results is also known to be important [4, 10]; in other words, results need to be organized into structured formats such as timelines, spreadsheets, and networks [4, 10]. We expect that schematizing may be even more criticial

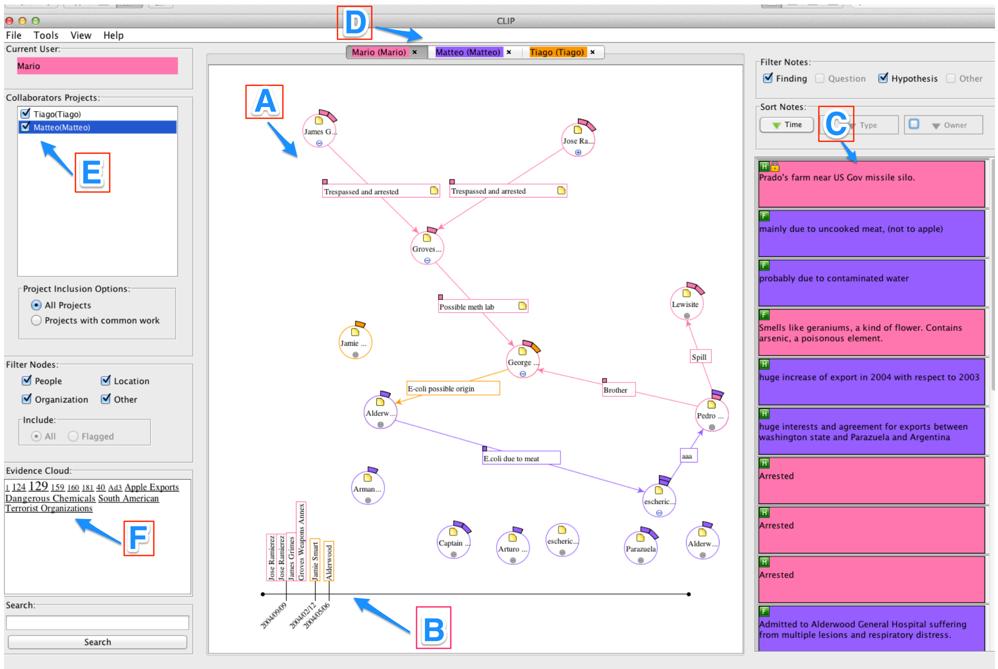

Figure 1: Screenshot of CLIP. A) Graph pane, to create a network diagram of people, locations, and events, B) Timeline, to see the timeline of events, C) List of notes, enabling easy review of all notes in one location, D) Tabs, to see collaborators' views, E) Merge option, to choose a collaborator's work to be merged with your own, F) Evidence cloud, to see the list of evidence and their frequencies.

for collaborative work, since the structure may additionally help with communication. CLIP therefore includes node-link graph and timeline schemas for representing findings.

2.2 Workspaces and Awareness

Providing both individual and group workspace enables collaborators to shift between tightly coupled and loosely coupled collaborative styles [6, 9, 10]. Various studies have found that collaboration styles directly affect outcomes (e.g. [10, 11, 16]). Teams that worked more closely were generally more successful; however, Kang and Stasko [11] found that collaboration was loose during information collection but tight when synthesizing findings. A good collaborative system, then, should encourage groups towards closer collaboration styles when there are relevant findings to be connected, but allow loose collaboration at other times. One way to encourage closer collaboration is through awareness mechanisms that provide information about collaborators' activities and findings. Some of these, such as the tabs in CoSpaces [13], place each person's information in a separate view. Users must then compare and reconcile different views, a potentially cumbersome process. People may also miss relevant changes that are hidden from view.

In contrast, integrating everyone's information into one view could cause disruption to individual work as the view constantly updates. One example system is Cambiera [9], which provided awareness cues about related searches conducted by a collaborator; these cues encouraged closer collaboration. We emphasize that Cambiera did not consider how collaborative brushing and linking could be applied to recorded notes and findings. Brennan et al. [2] did implement a visualization of externalizations and explored different ways to merge collaborators' content. Similarly, CoMotion [5] enabled analysts to simultaneously share views of data as well as notes. However, neither of these projects evaluated whether the merged or shared view was helpful to analysts in practice.

With CLIP, we address the need for individual workspace plus awareness of others' work by providing each user with flexible control over how much of the collaborators' information is shown in their view.

3. SYSTEM DESIGN

CLIP (Figure 1) is designed specifically to support synchronous collaborative intelligence analysis (e.g., solving a police mystery task). It facilitates analysis by providing a visual space for recording important entities found (e.g., suspicious people, locations, and events). Each recorded entity can be labeled with a type and attached to a note. To facilitate same-time collaboration, CLIP supports awareness through user-controlled sharing of work amongst team members. The following sections describe CLIP's features related to externalization and awareness support.

3.1 Externalization

CLIP provides a visual space for recording important entities and relationships (Figure 1A). Each recorded entity becomes a node in a node-link graph. The person who created each item is indicated by a unique colour.

Node-link graph: Each node represents an entity and has six main attributes: text, type, note, image, time-stamp and evidence list. *Text* appears in the node (e.g. a name like George Prado). *Type* assigns the entity to a category such as person, event, location, or organization. A *Note* can store extra textual information (e.g., "Who else knows George Prado?"). Based on its content, each note is labeled by the analyst as a hypothesis, question, finding, or other. Users can also set a note's scope to either public or private. *Image* is a picture that can be attached to a node. *Time-stamp* associates an entity with a specific date and time. *Evidence list* contains the name(s) of documents where the entity was found. Only the text and at least one evidence document are required for node creation; other node attributes are optional and can be updated any time.

Figure 2 shows how LCW appears on a node. Node text is placed in the middle of the circle. A yellow note icon indicates that a note is attached (Figure 2A). If the evidence list is not empty, segments are drawn on the outside of the circle, one for each evidence document attached to the node. These visual cues provide a quick overview of each node. By default, all notes are placed in the note panel (Figure 1C). Each note can be closed by the user later. When closed, the yellow note icon (inside the node within the graph) changes to red as a visual indication. Users can sort and filter notes in the note pane based on note type as well as time and owner. When a node with a Time-stamp attribute is created, the system automatically places a box with the node's text on the timeline (Figure 1B). Items on the timeline are ordered chronologically from left to right.

The implementation of CLIP supports full coordination of all views. When a node is enlarged to view details, the corresponding note and / or timeline item (if present) are highlighted by fading out other items, enabling the user to quickly identify related items. Similarly, selecting a note or timeline item highlights corresponding items in other views. Links represent relationships between captured entities. Each link has three main attributes: text, note and evidence list. These attributes mirror those of the nodes. Unlike a node's evidence list, a link's evidence list can be null.

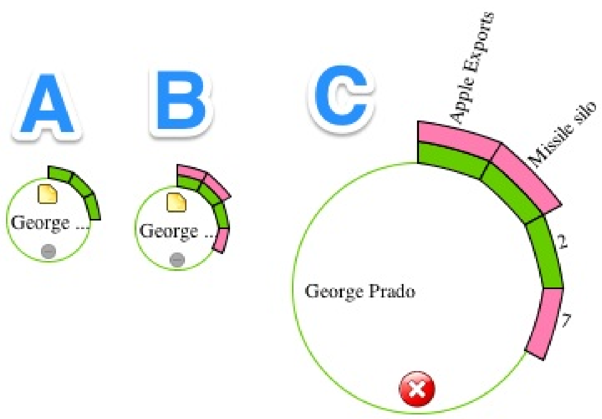

Figure 2: Node details. View of a node A) before and B) after a collaborator creates the same node. Colour coded segments reveal the common and different evidence. C) Enlarging a node shows details and also highlights related items in other views, such as notes, timeline items, and evidence items.

3.2 Awareness Support

To support awareness of collaborators' activities, instances of CLIP running on different machines communicate in real-time. Like Cambiera [9], distinctive colours are used to distinguish work by different people.

Partial merging: If another user has a node with the same name, then the local node is changed to notify the user that there is similar work. To keep changes in the local node subtle and yet noticeable, the only visual alteration is in the evidence list. Evidence lists of the local node and the collaborator's node are combined, and repeating evidence segments are stacked up. Figure 2 shows a node "George Prado" before and after a collaborator adds evidence. The colour of the local node is green (Figure 2B) and the colour of the collaborator is pink. In Figure 2B, the collaborators have two evidence items in common and other evidence identified by only one person or the other. In addition to changes in the common node, CLIP automatically combines all the notes together (ordered chronologically by default).Figure 2C shows an enlarged node to see more details about evidence.

Tabs: Each tab (Figure 1D) encompasses a view of the analysis work in another copy of CLIP. Tabs are labeled with the collaborator's unique colour and username. This enables fast recognition of who owns the work within a tab. Tabs show a node-link layout that is identical to the node-link layout created by the owner of that information.

Full merging: Figure 1 depicts an example of a fully merged view. The merged design enables the viewer to easily gain an understanding of how their collaborators' work relates to their own (e.g., what entities their collaborators are interested in and why, and what evidence they have found). Figure 1E is a list of collaborators' names that can be used to decide whose work to merge with your own. Checking the box next to a collaborator's name merges all of the collaborator's work into the local view. CLIP re-computes the graph layout and unites nodes with the same name. The primary user's layout is maintained as much as possible in order to preserve their mental map.

3.3 Implementation Details

Concurrently running instances of CLIP communicate using a peer-to-peer protocol. Changes by any collaborator are broadcasted across the group if a new node or link is created, or an existing node or link is updated or deleted. Upon receiving a message, the receiving end: 1) compares the local version of the collaborator's work (if existing) and updates the view accordingly. If a local view does not exist yet, a new tab is created that will encompass the collaborator's work. 2) compares the collaborator's content against the local content in search of common entities, which are then merged in the display. In this version, we consider two entities related if the nodes have the same name.

4. DISCUSSION AND FUTURE WORK

Our research extends earlier work on thinking spaces by integrating LCW features that reveal relationships between collaborators' findings. Results of our user study (not discussed here because of space limitations) showed that LCW clearly supported groups in this collaborative sensemaking task. CLIP groups achieved significantly better scores, coordinated and communicated more effectively, and reported that they depended on LCW to maintain awareness of each others' work. CLIP groups were able to focus their oral communication on coordination activities rather than using oral communication as a (disruptive) awareness channel.

CLIP is a prototype tool implemented specifically to explore the benefits of LCW. While it was an effective thinking space for the intelligence analysis task in our study, additional work would be needed to extend it for more general use. To begin with, CLIP aims to support the sensemaking loop, and therefore provides no support for information foraging. Combining CLIP with a complementary tool like Cambiera [9] may be an effective way to support both phases. We strongly suspect information foraging support would be crucial when dealing with a larger document set. Another limitation of CLIP is that it does not automatically identify entities or relationships between documents; we extracted and encoded these manually as a demonstration. Many algorithms exist for finding these relationships automatically. For this reason, we would like to integrate CLIP's thinking space and LCW features into a document analytics tool such as Jigsaw [15] that automatically extracts entities and relationships from the documents.

It is also not clear how well CLIP's visual representation of the thinking space will scale to very large and complex problems or data sets other than document collections. Different visual representations could be required to support different types of data and analytical questions. The structures (e.g., node-link graph) used to organize items may not scale well even for individuals, and with multiple analysts, keeping track of collaborators' changes and updates to such a large representation may be impossible.

5. CONCLUSION

CLIP demonstrates how linked common work can be employed within 'thinking spaces' to support sensemaking during collaborative analytics. LCW can significantly improve analytic outcomes at a collaborative intelligence task.

6. ACKNOWLEDGEMENTS

We wish to thank members of the VisID group at the University of Victoria for their helpful feedback and suggestions. Also many thanks to Rock Leung and our other colleagues at SAP for their feedback and suggestions on system design. This research was funded by SAP, NSERC, and GRAND.

7. REFERENCES

[1] E. A. Bier, E. W. Ishak, and E. Chi. Entity workspace: an evidence file that aids memory, inference, and reading. In *Intelligence and Security Informatics*, pages 466–472. Springer, 2006.

[2] S. E. Brennan, K. Mueller, G. Zelinsky, I. Ramakrishnan, D. S. Warren, and A. Kaufman. Toward a multi-analyst, collaborative framework for visual analytics. In *Proc. Visual Analytics Science And Technology (VAST)*, pages 129–136. IEEE, 2006.

[3] Y. Chen, J. Alsakran, S. Barlowe, J. Yang, and Y. Zhao. Supporting effective common ground construction in asynchronous collaborative visual analytics. In *Proc. Visual Analytics Science and Technology (VAST), 2011 IEEE Conference on*, pages 101–110. IEEE, 2011.

[4] G. Chin Jr, O. A. Kuchar, and K. E. Wolf. Exploring the analytical processes of intelligence analysts. In *Proc. SIGCHI Conf. on Human Factors in Computing Systems*, pages 11–20. ACM, 2009.

[5] M. C. Chuah and S. F. Roth. Visualizing common ground. In *Proc. Int. Conf. on Information Visualization*, pages 365–372. IEEE, 2003.

[6] C. Gutwin and S. Greenberg. Design for individuals, design for groups: tradeoffs between power and workspace awareness. In *Proc. ACM Conf. on Computer supported cooperative work*, pages 207–216. ACM, 1998.

[7] J. Heer, F. van Ham, S. Carpendale, C. Weaver, and P. Isenberg. Creation and collaboration: Engaging new audiences for information visualization. In *Information Visualization*, pages 92–133. Springer, 2008.

[8] M. S. Hossain, C. Andrews, N. Ramakrishnan, and C. North. Helping intelligence analysts make connections. In *Scalable Integration of Analytics and Visualization*, 2011.

[9] P. Isenberg and D. Fisher. Collaborative brushing and linking for co-located visual analytics of document collections. In *Computer Graphics Forum*, volume 28, pages 1031–1038. Wiley Online Library, 2009.

[10] P. Isenberg, D. Fisher, S. A. Paul, M. R. Morris, K. Inkpen, and M. Czerwinski. Co-located collaborative visual analytics around a tabletop display. *Visualization and Computer Graphics, IEEE Transactions on*, 18(5):689–702, 2012.

[11] Y.-a. Kang and J. Stasko. Characterizing the intelligence analysis process: Informing visual analytics design through a longitudinal field study. In *Proc. Visual Analytics Science and Technology (VAST)*, pages 21–30. IEEE, 2011.

[12] N. Mahyar, A. Sarvghad, and M. Tory. Note-taking in co-located collaborative visual analytics: Analysis of an observational study. *Information Visualization*, 11(3):190–204, 2012.

[13] N. Mahyar, A. Sarvghad, M. Tory, and T. Weeres. Observations of record-keeping in co-located collaborative analysis. In *Proc. 46th Hawaii Int. Conf. on System Sciences (HICSS)*, pages 460–469. IEEE, 2013.

[14] P. Pirolli and S. Card. The sensemaking process and leverage points for analyst technology as identified through cognitive task analysis. In *Proc. Int. Conf. on Intelligence Analysis*, volume 5, pages 2–4, 2005.

[15] J. Stasko, C. Görg, and Z. Liu. Jigsaw: supporting investigative analysis through interactive visualization. *Information visualization*, 7(2):118–132, 2008.

[16] K. Vogt, L. Bradel, C. Andrews, C. North, A. Endert, and D. Hutchings. Co-located collaborative sensemaking on a large high-resolution display with multiple input devices. In *Proc. Human-Computer Interaction–INTERACT 2011*, pages 589–604. Springer, 2011.

[17] W. Willett, J. Heer, J. Hellerstein, and M. Agrawala. Commentspace: structured support for collaborative visual analysis. In *Proc. SIGCHI Conf. on Human Factors in Computing Systems*, pages 3131–3140. ACM, 2011.

[18] W. Wright, D. Schroh, P. Proulx, A. Skaburskis, and B. Cort. The sandbox for analysis: concepts and methods. In *Proc. SIGCHI Conf. on Human Factors in computing systems*, pages 801–810. ACM, 2006.