Supporting Sensemaking during Collocated Collaborative Visual Analytics

by

Narges Mahyar
B.Sc., Azad University of Tehran, 1998
M.Sc., University of Malaya, 2008

A Dissertation Submitted in Partial Fulfillment of the
Requirements for the Degree of

DOCTOR OF PHILOSOPHY

in the Department of Computer Science

Supporting Sensemaking during Collocated Collaborative Visual Analytics

by

Narges Mahyar
B.Sc., Azad University of Tehran, 1998
M.Sc., University of Malaya, 2008

Supervisory Committee

Dr. Melanie Tory, Supervisor
(Department of Computer Science)

Dr. Margaret-Anne Storey, Departmental Member
(Department of Computer Science)

Dr. Adel Guitouni, Outside Member
(Peter B. Gustavson School of Business)

Supervisory Committee

Dr. Melanie Tory, Supervisor
(Department of Computer Science)

Dr. Margaret-Anne Storey, Departmental Member
(Department of Computer Science)

Dr. Adel Guitouni, Outside Member
(Peter B. Gustavson School of Business)

ABSTRACT

Sensemaking (i.e. the process of deriving meaning from complex information to make decisions) is often cited as an important and challenging activity for collaborative technology. A key element to the success of collaborative sensemaking is effective coordination and communication within the team. It requires team members to divide the task load, communicate findings and discuss the results. Sensemaking is one of the human activities involved in visual analytics (i.e. the science of analytical reasoning facilitated by interactive visual interfaces). The inherent complexity of the sensemaking process imposes many challenges for designers.

Therefore, providing effective tool support for collaborative sensemaking is a multifaceted and complex problem. Such tools should provide support for visualization as well as communication and coordination. Analysts need to organize their findings, hypotheses, and evidence, share that information with their collaborators, and coordinate work activities amongst members of the team. Sharing externalizations (i.e. any information related to the course of analysis such as insights, hypotheses, to-do lists, reminders, etc recorded in the form of note/ annotation) could increase awareness and assist team members to better communicate and coordinate their work activities. However, we currently know very little about how to provide tool support for this sort of sharing.

This thesis is structured around three major phases. It consists of a series of studies to better understand collaborative Visual Analytics (VA) processes and challenges, and empirically evaluate design ideas for supporting collaborative sensemaking. I investigate how collaborative sensemaking can be supported during visual analytics by a small team of collocated analysts. In the first phase of this research, I conducted an observational study to better understand the process of sensemaking during collaborative visual analytics as well as identify challenges and further requirements. This study enabled me to develop a deeper understanding of the collocated collaborative visual analytics process and activities involved. I found that record-keeping plays a critical role in the overall process of collaborative visual analytics. Record-keeping involves recording any information related to the analysis task including visualization snapshots, system states, notes, annotations and any other material for further analysis such as reminders and to-do lists. Based on my observations, I proposed a characterization of activities during collaborative visual analytics that encompasses record-keeping as one of the main activities. In addition, I characterized notes according to their content, scope, and usage, and described how they fit into a process of collaborative data analysis. Then, I derived guidelines to improve the design of record-keeping functionality for collocated collaborative visual analytics tools.

One of the main design implications of my observational study was to integrate record-keeping functionality into a collaborative visual analytics tool. In order to examine how this feature should be integrated with current VA tools, in the second phase of this research, I designed, developed and evaluated a tool, CoSpaces (Collaborative Spaces), tailor-made for collocated collaborative data analysis on large interactive surfaces. Based on the result of a user study with this tool, I characterized users' actions on visual record-keeping as well as their key intentions for each action. In addition, I proposed further design guidelines such as providing various views of recorded material, showing manually saved rather than automatically saved items by default, enabling people to review collaborators' work unobtrusively, and automatically recommending items related to a user's analytical task.

In the third phase, I took supporting record-keeping activities in the context of collaborative sensemaking a step further to investigate how this support should be designed to facilitate collaboration. To this end, I explored how automatic discovery and linking of common work can be employed within a "collaborative thinking space" (i.e. a space to enable analysts to record and organize findings, evidence, and hypotheses, also facilitate the process of sharing findings amongst collaborators), to facilitate synchronous collaborative sensemaking activities in visual analytics. The main goal of this phase was to provide an environment for analysts to record, organize, share and connect externalizations. I expected that this would increase awareness among team members and in turn would enhance communication and coordination of activities. I designed, implemented and evaluated a new tool, CLIP (Collaborative Intelligence Pad), that extends earlier thinking spaces by integrating new features that reveal relationships between collaborators' findings. Comparing CLIP versus a baseline tool demonstrated that linking collaborators' work led to significant improvement in analytical outcomes at a collaborative intelligence task. Groups using CLIP were also able to more effectively coordinate their work, and held more discussion of their findings and hypotheses. Based on this study, I proposed design guidelines collaborative VA tools.

In summary, I contribute an understanding for how analysts use VA tools during collocated collaboration. Through a series of observational user studies, I investigated how we can better support this complex process. More specifically, I empirically studied recording and sharing of analytical results. For this purpose, I implemented and evaluated two systems to be able to understand the effects of these tools on collaboration mechanics. These user studies along with various literature surveys on each specific topic resulted in a collection of guidelines for supporting and sharing externalizations. In addition, I proposed and evaluated several mechanisms to increase awareness among team members, resulting in more effective coordination and communication during the collaborative sensemaking process. The most novel contributions of this research are the identification and subsequent characterization of note taking behaviours as an important component of visual data exploration and analysis. Moreover, the design and evaluation of CLIP, providing preliminary evidence in support of automatically identifying and presenting relationships between collaborators' findings.

Contents

List of Tables

List of Figures

ACKNOWLEDGEMENTS

I would like to express my deepest appreciation to the people who inspired and supported me during the completion of my dissertation, including those not mentioned here by name: there were so many that it is impossible to thank everyone individually.

First, I offer heartfelt gratitude to my supervisor, Dr. Melanie Tory. This thesis would not have been possible without her continuous guidance and support, as well as her pushing me to think about things that seemed impossible. I also value the excellent example she sets as a successful women and professor.

Thanks also go to members of my supervisory committee: Dr. Margaret-Anne Storey, Dr. Adel Guitouni and external examiner Dr. Jeffrey Michael Heer. Their constructive criticism along with their unique and important perspectives helped me to improve the quality of this research noticeably.

My colleagues at SAP, especially Rock Leung and Michael McAllister, provided invaluable support through four years of my PhD research in the SAP ARC project, for which I am profoundly grateful. Also thanks to my colleagues at VisID lab for their constant feedback on my research and their participation in many pilot studies. Thanks to Leandro Collares and Wanda Boyer for their assistance with data analysis of the third study.

I also appreciate SAP, NSERC and the GRAND research networks for funding me during my doctoral studies.

Thank you to my family and friends. Deepest appreciation goes to my parents for their love, support and encouraging me to make my own decisions. Thanks to my sisters for their love, patience, and unwavering belief in me, and to my brothers for showing me the value of hard work and being persistent. To me that is what PhD stands for: Persistence, hard work, Determination. Thank you to my son, Iliya, for being a great motivation and inspiration to complete this work and for teaching me through his endless joyful spirit how to enjoy each and every moment. Last, but certainly not least, thank you to my loving, supportive, and patient husband, Ali. I could not have enjoyed this process without your faithful support and encouragement.

To my husband Ali,

For his calm spirit that makes anything possible,
his endless love and kindness, and
amazing support and patience.

And to my son, Iliya, who reminds me of "The little prince",

Not for his long curly hair,
not even for his big curious eyes.
But for his philosophical spirit,
his limitless appetite to learn, and his endless questions!

PUBLICATIONS

The materials presented in this thesis have been previously published in different venues. After each reference, I refer to chapters that present the material.

Journal Articles

- **Narges Mahyar**, and Melanie Tory, Supporting Communication and Coordination in Collaborative Sensemaking, IEEE Transactions on Visualization and Computer Graphics. (to appear, pages:10).

  Material from this publication appears in chapter 5.

- **Narges Mahyar**, Ali Sarvghad, and Melanie Tory, Note Taking in Co-located Collaborative Visual analytics: Analysis of an Observational Study, Information Visualization, vol. 11, no. 3, pp. 190-204, July 2012.

  Material from this publication appears in chapter 3.

Conference Papers

- **Narges Mahyar**, Ali Sarvghad, Melanie Tory, and Tyler Weeres, Observations of Record-Keeping in Co-located Collaborative Analysis, HICSS 2013, pp. 460-469, Jan. 2013.

  Material from this publication appears in chapter 4.

- N**arges Mahyar**, Ali Sarvghad, and Melanie Tory, A closer look at note taking in the co-located collaborative visual analytics process, IEEE Visual Analytics Science and Technology (VAST10), pp. 171-178, 2010. [Selected for publication in the "Information Visualization" journal].

  Material from this publication appears in chapter 3.

Workshop Papers

- **Narges Mahyar**, and Melanie Tory, CLIP: A visual thinking space to support collaborative sensemaking and reasoning, Graphics, Animation and New Media (GRAND) NCE AGM, 2014. [Best honorable mention research note].

  Material from this publication appears in chapter 5.

- **Narges Mahyar**, Ali Sarvghad, Melanie Tory and Tyler Weeres CoSpaces: Workspaces to Support Co-located Collaborative Visual Analytics, DEXIS 2011, Nov 2011.

  Material from this publication appears in chapter 4.

- **Narges Mahyar**, Ali Sarvghad, and Melanie Tory, "Roles of Notes in Co-located Collaborative Visualization," Workshop on Collaborative Visualization on Interactive Surfaces (CoVis 2009), Oct. 2009.

  Material from this publication appears in chapter 3.

- Ali Sarvghad, **Narges Mahyar**, and Melanie Tory, "History Tools for Collaborative Visualization, " Workshop on Collaborative Visualization on Interactive Surfaces (CoVis 2009), Oct. 2009.

  Material from this publication appears in chapters 3 and 4.

Chapter 1

Introduction

In many disciplines people need to analyze and make sense of large and complex datasets. Due to the size and complexity of data, often analysts work together to solve complex problems and make decisions. In addition, many domain problems are interdisciplinary in nature and require analysts with different expertise to work together. By collaborating, analysts with different perspectives can each contribute their own expertise to improve the quality of the solutions. In order to analyze such complex datasets, experts often employ various analytical techniques and rely on tools that support collaboration. In particular, Visual Analytics (VA) facilitates exploring and understanding complex data through visualization and analytical techniques. According to Cook and Thomas [18] VA is "the science of analytical reasoning facilitated by interactive visual interfaces". Sensemaking is one of the human activities involved in VA . Sensemaking is a creative and retrospective process of deriving meaning from information to help in making decisions [113]. Similar to any creative process, it is not structured and requires interaction with the data to discover interesting facts [27]. Sensemaking requires iteratively detecting and adjusting patterns; the analyst iteratively creates hypotheses and validates them through various interactions.

Although there is a rich body of research that investigates groupware and VA tools, there is relatively less research that investigates design issues and requirements specific to collaborative sensemaking in VA. The inherent complexity of the sensemaking process imposes many challenges for designers. According to the Pirolli and Card model [84], sensemaking is comprised of two top level loops, the information foraging loop and the sensemaking loop (See Chapter 2 for more details). The focus of the information foraging loop is discovery of information within the data space and the objective of the sensemaking loop is relating the findings and building an overall understanding of the data to solve the task. The main focus of prior research on collaborative sensemaking has been on supporting the information foraging loop.

However supporting the sensemaking loop, especially in collaborative situations, imposes a lot of complexities for designers of VA tools. While this is really important for many collaborative real life problems, we still need to gain a better understanding of how the collaborative sensemaking cycle happens in a natural setting. In addition, there has been little investigation of how collaborative VA tools can provide better support for analysts' needs during the sensemaking loop. In this spirit, Cook and Thomas [18] name the design of collaborative visualization tools as a grand challenge for visualization research. In a more recent survey in the field of collaborative visualization, Isenberg et al. [44] still refer to this challenge. One particular challenge that they mention is the difficulty of keeping all the relevant entities clear in one's mind, to remember the context in which they were discussed, and to connect them to other activities that were noted. In the collaborative sensemaking process, analysts need to deal with both their own and others' past discoveries. Prior VA research also mentioned that analysts heavily rely on insights and findings discovered in the course of analysis [37,43,49,61,97]. Therefore, the ability to record findings and relate them to each other has been suggested to improve the sensemaking process. For instance, Isenberg et al. [46] suggested that sharing of analytical findings within a group could result in closer collaboration and assist teams with decision making and problem solving. Therefore, providing support for recording, schematizing and sharing findings, hypotheses and questions within the group seems to be essential. However, there are many interesting questions related to this research direction. For instance: How do collaborating analysts record and share their findings without support tools? How can we better design collaborative tools to support this process?

All of the challenges in supporting the collaborative sensemaking process, along with the critical need to support collaborative sensemaking activities in many real life situations, motivated me to focus on this research area (See Chapter 2 for more details about the challenges and real life scenarios). Particularly, I chose to focus on investigating how to provide support for recording and sharing analytical results in a collaborative setting in collocated scenarios. In the remaining sections, I elaborate on the research problems, scope, methods, and contributions of this research.

1.1 Thesis Problem

In this thesis, I investigate how the collaborative sensemaking process happens in a collocated setting and how collaborative VA tools can provide better support for analysts' needs during the sensemaking process. More specifically, I investigate how to design tools that provide better support for collaborative sensemaking during collocated visual analytics.

The goal of this thesis is to develop a theoretical understanding of how tool support should be designed to better facilitate the collaborative sensemaking process in visual analytics. The results of this research are expected to inform the design of future collaborative VA systems.

I started this research by conducting an initial literature survey which revealed many challenges for designing collaborative VA tools. However, it was still not quite clear how people work together using a VA tool or which challenges could be more important during this process. I was interested to explore and understand the specific challenges for collocated collaborative VA. Therefore, I first conducted an observational study to gain a better understanding of the process of collaborative visual analytics and open issues in this domain. In chapter 3, I discuss the first research question:

RQ1: How can collaborative record-keeping activities be characterized in the process of visual analytics for a small team of collocated collaborators?
This research question emerged during an exploratory study aiming to explore challenges in collocated VA. The results clearly indicated the importance of recording and sharing results during collocated collaborative VA. Therefore, I characterized record-keeping activities (i.e. recording past visualization states as well as externalizations in the form of notes and annotations) as part of the collaborative analysis process. Then, I suggested design guidelines to better support record-keeping activities in collaborative VA. The results of this study led to defining the next research question:

RQ2: How can record-keeping functionality be integrated into a collaborative VA tool in a way that supports collaborative mechanics?
The importance of record-keeping activities was one of the main design implications of the first phase. Therefore, in the second phase I investigated how to integrate record-keeping functionality into a VA tool (Chapter 4). At the time, there were

no collaborative VA tools designed for tabletops that supported this functionality. Therefore, I designed and developed CoSpaces (Collaborative Spaces) to support record-keeping activities on large interactive surfaces. The main design objective of this prototype was providing a mechanism for recording and sharing of results.

The results of evaluating CoSpaces showed that providing support for record-keeping and unobtrusive sharing of results facilitated the VA process. However, I found that it was too cumbersome for people to review each other's work through separate workspaces. This motivated me to investigate other ways to increase awareness[1] [2] among team members. I speculated that integrating collaborators' work more directly so that they would automatically see linkages between their own and others' work would be a key element for better communication and coordination. This led me to explore how I could link collaborators' work visually and what effects this would have on the collaborative analysis process:

RQ3: What are the effects of linking collaborators' related externalizations to one another on awareness, performance, communication and coordination?

To answer this question, I started by conducting an in depth literature survey to understand requirements for the collaborative sensemaking process. The literature review revealed a lot of challenges to support collaborative sensemaking. To better understand users' needs and challenges I looked into empirical studies in the field of collaborative sensemaking. The result of this survey is a list of design guidelines discussed in Chapter 2 section 2.5. Considering these guidelines as well as results of my previous study, I proposed a collaborative thinking space that integrated linking of common work in order to increase awareness among team members. To explore the effects of discovering and linking common work, I designed and developed, CLIP (Collaborative Intelligence Pad). The core feature of CLIP was increasing awareness with employing LCW (Linked Common Work) technique.

---

[1] In CSCW literature awareness has been defined in different ways. For instance Gutwin et al. [32] define awareness for groupware as group structural awareness, social awareness, informal awareness and workspace awareness. Although in collocated collaboration people have the advantage of being co-present, still this kind of awareness could be lost particularly while shifting from loosely to tightly coupled collaboration and vice versa [33].

[2] In this thesis I refer to awareness in a general manner as "up to moment understanding of each other's work, including both activities that people are working on as well as results and evidence that they have found'.'

1.2 Thesis Scope

This thesis focuses on supporting visual analytics in collocated scenarios where a small team of collaborators works together synchronously to make sense of data. I chose collocated situations for many reasons. In Chapter 2 (section 2.1.1) I discuss the benefits that collaboration offers in many disciplines, as well as the critical need to support collocated collaborative scenarios in real life. In addition, collocated situations may represent the best case for collaboration, as users have all of the advantages of working synchronously together at the same place. This is important because problems that occur in collocated situations may be likely to occur in other types of collaboration as well. Furthermore, in collocated situations, researchers can more easily examine how team members collaborate in real-time. Direct interaction with all the group members is a great opportunity to understand their needs and challenges and it is much easier to conduct post study interviews in collocated studies.

The workspace organization in this research includes both single shared workspace (large interactive surfaces), and multiple displays. For the first two phases of this research I chose large interactive displays due to their prominent use in collocated collaborative scenarios. Large interactive displays have great potential for groupware applications and have been used in many domains such as science, engineering, intelligence analysis, and health care [88]. In particular, large interactive displays have been used in collaborative sensemaking (e.g. [43, 46, 75, 104, 109]). While they might have been used successfully in other contexts, there are challenges involved in using these displays for visual analytics, especially in regards to note taking. According to [33, 110] the tradeoff of using shared displays is that they can compromise individual exploration of the problem space. In addition, in phase two of this research (Chapter 4), I observed difficulties in taking notes on a large interactive surface using an on-screen keyboard. However, note taking and externalizations are the main focus of this research. Moreover, due to many reasons, a lot of real life collaborative scenarios still happen in traditional settings (multiple displays). In many collocated VA sessions, analysts bring their own laptops or use their own desktops. Therefore, in the third phase, I used multiple desktops to encourage equity of participation and ensure the ease of note taking.

Application domains investigated in this research are the business domain and intelligence analysis. With a different focus, both domains include a set of systematic and analytical processes in order to explore complex datasets, find trends, patterns and correlations to solve a problem. The application areas are not the main focus of the thesis but rather the main focus is on understanding the collaborative VA process to inform the VA community in general.

Several motivations led to choosing these domains. First, collocated collaboration is a frequent use case for both domains [18, 21, 46, 55, 83, 84, 109]. In both fields, analysts often work together to solve complex problems and make decisions, particularly when each user has unique expertise to offer. Another motivation for considering the business domain in this research was close collaboration with SAP (the industrial partner that partially funded this research).

It should be noted that both fields have been investigated broadly and numerous visualization tools exist in these domains. However, most of the previously existing tools were designed for single users. This research aims to establish a foundation for extending those tools to support collaborative work.

1.3 Methodological Approach

For each phase of the research, I paid special attention and consulted many useful references to choose appropriate methods (e.g. [19, 20, 70, 105]). Qualitative methods[3] are common practice for understanding team dynamics and individuals' behaviour in the collaborative work domain (e.g. [31, 46, 48, 77, 94, 102, 106]). However to avoid researcher bias, it is important to follow a structured and established approach. One well known approach to offset some of the flaws of each methodological approach (qualitative vs. quantitative) is using mixed methods [19, 20, 70].

For all phases of this research I chose a controlled user study method. User studies are becoming standard practice in visualization research and design, as a way of both understanding users and evaluating visualization tools [105]. Notably, the user study is not one single research method; rather it encompasses many different empirical methods involving observations of participants' behaviour. All of these studies began with a series of pilot studies to offset some possible flaws.

In addition to many considerations for choosing a method (such as how well it fits the research question and how to minimize the effect of its flaws), its practicality is an important factor to consider [20]. Observing and understanding collaborative situations, especially around large interactive surfaces, has its own challenges. For

[3]More detailed descriptions of qualitative methods and details on the coding process can be found in Creswell et al. [19].

instance, most of the time field or ethnography studies are not possible due to the lack of required hardware and accessibility to workplaces using this hardware. Other key challenges involved in employing ethnographic field studies are finding the time and resources needed to collect and analyze potentially large sets of data, and establishing a relationship with the studied organization [20]. The use of large interactive surfaces for the first two phases of this research made it almost impossible to choose this approach. For this research, a combination of qualitative and quantitative approaches seemed to suit this research the best. Qualitative observation of behviours was necessary to understand groups' challenges and design requirements. Quantitative methods, on the other hand, was complementary (e.g. measuring variables such as time, and accuracy).

In the first phase, I used an exploratory qualitative approach, because the study was exploratory rather than aiming to confirm specific questions. I conducted an abstract lab study but then analyzed the results using a qualitative coding approach. I was particularly interested to study people collaboratively analyzing information with a single user support tool around large interactive surfaces. The goal of this research was to understand how to design support tools for collocated collaborative VA on large interactive surfaces. I gathered data from different sources including the experimenter's observations, captured video, audio and interview materials. The first plausible step to gain a better understanding of the process and users' needs was to study this process without using software support tools to eliminate the possible effects of tool used. This step was taken by previous researchers (e.g. [48, 86, 94]). Therefore, the natural next step seemed to be studying group behaviour around a large interactive surface to understand people's challenges and possible needs before designing a new tool. The third step to this end is to prototype tools based on guidelines that emerge from such studies in order to validate them better.

In phase two, while an alternative method could be a usability study, again I chose a controlled user study with a mixed methods approach. I was mainly interested to observe users' behaviour and their collaboration mechanics while they used record-keeping functionality within a VA tool. However, because it was a prototype tool, it was critical to understand users' possible challenges during the use of the prototype. For this reason, I used a questionnaire to receive user' feedback on the new features that I had designed. Also I eliminated assessing the usability of the tool as a whole and focused on evaluating the new design decisions that I had employed in this prototype. Therefore, a combination of methods including observational techniques, questionnaires, system logs, and interviews was used to assess how the prototype tool was used, users' challenges during tool use, and how the tool affected their collaboration. Captured videos were coded using a carefully developed coding scheme. To reduce the interpretation problem and experimenter bias, two independent coders observed and coded the qualitative data. We coded data independently and then shared our results. In order to come to an agreement, sometimes we had to conduct another round of analysis to gather enough evidence.

In the third phase, similar methods were used (i.e. controlled user study). Similar to the previous study, I evaluated the tool by focusing on the new aspects that had been added to this system compared to the previous tools. Post study questionnaires were designed to gather users' feedback on specific features of this tool that were relevant to my hypotheses. In addition, any problems related to the use of the tool during the sessions were recorded by the experimenter. However, in addition to a questionnaire, an in-depth quantitative data analysis using two independent coders was done for analyzing conversations between collaborators. The coding scheme was carefully designed to measure data related to my hypotheses, specifically including measures of communication effectiveness, coordination, and awareness. Using an iteratively built coding scheme, I categorized each instance of conversation. In addition to the quantitative analysis of conversations, both observations made by the experimenter and recorded videos, logs and screen shots were qualitatively analyzed to better understand the effects of the design ideas on collaboration. For this study, I used a standard task and dataset with a ground truth that enabled me to quantitatively evaluate the analytical outcome score. I employed a between subject comparison to a base tool to understand the effects of the new design decisions on analysis outcomes and collaboration mechanics. For this purpose I used metrics such as the analytical outcome score, frequency of referring to the visualization tool and different codes to understand the communication add coordination. In Chapter 5 I discuss the coding in detail as well as how I measured different aspects of the collaborative process.

1.4 Thesis Contributions

The main contributions of the first phase of this research, as described in Chapter 3, are as follows:

C1: A characterization of the collocated visual analytics process and ac-

tivities during collocated VA, demonstrating the importance of record-keeping activities during this process.
I developed a characterization of the collocated collaborative visual analytics process and activities. I discussed the importance of record-keeping and how record-keeping activities fit into the overall process of collaborative VA.

C2: A categorization of notes based on their contents, scope, and usage.
I proposed a categorization of notes based on notes taken by participants during the visual analytics tasks. This rich understanding of note categories and usage could be beneficial for developers of collaborative VA tools.

C3: A set of guidelines to improve the design of record-keeping for collocated collaborative visual analytics tools.
I derived a set of guidelines to improve the design of record-keeping and note taking during collaborative VA.

In the second phase, in order to propose more practical design considerations, I designed and implemented a prototype tool for collaborative visual analytics on tabletops. Major contributions of the second phase, as described in Chapter 4, are below:

C4: Validated the importance of externalization support during collaborative VA and revealed new ways to provide externalization support.
This study corroborated the value of record-keeping for supporting collaboration during visual analytics process. This study extended the importance of record-keeping support to the collocated collaborative settings.

C5: A Characterization of users' actions on visual record-keeping as well as their key intentions for each action.
Actions and intentions and also their dependence on the analytical phase and collaboration style derived from this study. This characterization can be used as a fundamental base to improve design of collaborative VA tools.

C6: Design guidelines for collaborative VA tools to improve record-keeping support.

These guidelines include further suggestions to improve record-keeping support.

In the third phase of this thesis, I designed and implemented a collaborative thinking space that automatically discovered and notified collaborators of recorded common work. Main contributions of this phase, as described in Chapter 5, are listed as below:

C7: Introduction, implementation and evaluation of LCW techniques in the context of collaborative sensemaking.
I explored how automatic discovery and linking of common work can be employed to facilitate collocated collaborative sensemaking activities in visual analytics. The main goal of this phase was to provide an environment for analysts to record, organize, share and connect externalizations. I designed, implemented and evaluated a new tool, CLIP, that integrates features that reveal relationships between collaborators' findings. Results of this study demonstrated that linking collaborators' work led to significant improvement in analytical outcomes, improved communication and coordination and increased awareness at a collaborative intelligence task.

C8: Design guidelines to embed collaborative thinking spaces into collaborative VA tools.
Based on this study, I proposed different ways to improve awareness, help with coordination and communication.

C9: A set of metrics to measure awareness, coordination and communication for collaborative VA. Another contribution of this phase include an in depth quantitative evaluation method and metrics for measuring coordination, communication and awareness during a collaborative VA session.

1.5 Thesis Outline

This thesis is structured around the three main research questions and corresponding phases. It consists of a series of studies to better understand collaborative VA processes and challenges, and empirically test and evaluate design ideas in order to inform the VA community. The next chapter describes , chapters 3 to 5 describe the three phases of this research, and chapter 6 discusses lessons learned and future directions. In chapter 7, I conclude by summarizing the thesis and addressing the contributions of this research.

Chapter 2 Related Work

I introduce relevant background material related to collaborative visualization, the collaborative VA process, record-keeping and collaborative sensemaking. This chapter presents related work relevant to this thesis.

Chapter 3 Understanding Collocated Collaborative Visual Analytics Processes and the Role of Record-keeping

RQ1, C1, C2, and C3

I present a rich understanding of collaborative VA processes and challenges derived from an observational user study. In addition, I propose a characterization of the collaborative VA process along with a categorization of notes taken during the study. Then, I discuss design suggestions to support record-keeping in VA tools.

Chapter 4 CoSpaces - A System to Support Record-Keeping in Collocated Collaborative Visual Analytics

RQ2, C4, C5, and C6

I report on a user study in which I examined a prototype tool that incorporates record-keeping for collaborative VA. I derive a list of actions and intentions in regard to the use of the visual record-keeping module and discuss their relevance to the collaboration phases. I also further discuss design requirements to integrate record-keeping into VA tools.

Chapter 5 CLIP - A System to Support Communication and Coordination in Collaborative Sensemaking

RQ3, C7, C8 and C9

I designed, implemented and evaluated a new tool, CLIP, that extends earlier thinking spaces by revealing relationships between collaborators' findings. I demonstrated that CLIP significantly improved analytic outcomes at a collaborative intelligence task (in comparison to baseline groups). CLIP groups were able to more effectively coordinate their work, and held more discussion of their findings and hypotheses. Based on my results, I propose design guidelines for embedding collaborative thinking spaces into collaborative VA tools.

Chapter 6 Discussion and Future Work
I discuss lessons learned, limitations, threats to validity and future directions.

Chapter 7 Summary and Contributions
I summarize the thesis and address the contributions.

Chapter 2

Related Work

This thesis is focused on supporting collaborative sensemaking during the visual analytics process and addressing challenges for a small team of collocated collaborators during this process. Collaborative sensemaking occurs in many domains, including healthcare, science, emergency services, business and intelligence analysis (e.g. [82, 112]). While many previous guidelines related to designing single user VA tools still apply to collaborative VA, supporting collaborative sensemaking is still a challenge. Interfaces, visualizations, and interaction techniques should be designed to specifically address the needs and requirements of multiple collaborators.

To provide better support for collaborative problem solving, it is important to understand how analysts tackle complex tasks such as intelligence analysis or business problems. I began with a brief review of collaborative visualization in general and more specific to this thesis, collocated scenarios and their occurrence in real life situations. Then I review visual analytics literature in a nutshell and the analytic process employed by people during collaborative VA tasks. Since supporting record-keeping is the main theme of this research, I elaborate on its role and importance as well as current support for different forms of record-keeping in visual analytics. Then I discuss design considerations and challenges more specific to collocated collaborative VA. Finally, I define sensemaking and describe the sensemaking process as it is currently understood.

2.1 Collaborative Visualization

Collaborative visualization is a relatively new discipline that takes advantage of collaboration and visualization to derive information and insight from data. A survey of the current literature on collaborative visualization [44] produced a varying number of definitions. In this thesis, I will adopt the current definition introduced by Isenberg et al. [44]. Based on their extensive review of the prior research in this field, they define collaborative visualization as “the shared use of computer-supported, (interactive,) visual representations of data by more than one person with the common goal of contribution to joint information processing activities.”

This discipline exploits many advantages offered by collaboration and visualization. For instance, sometimes extracting relevant information from the flood of data is simply too complex for an individual. In collaborative projects, the task load can be shared among individuals on a team [18]. This allows team members to each focus on different tasks or subtasks. In addition, interpretations of the dataset and validity of solutions by a group is less susceptible to individual biases and errors that might occur in analysis by an individual. Collaboration can improve the analytical outcome by allowing analysts to discuss, evaluate and validate different hypotheses [36,44,48]. Visualization provides abstract representation of the underlying data that enables analysts to visually discover trends, patterns, outliers, etc. Although providing new and exciting benefits, this field also introduces new and unique challenges and obstacles to tackle. In the following sections I address some issues in regards to collaborative visual analytics and design for collocated collaboration around large interactive displays.

2.1.1 Collocated Collaboration Scenarios

Collaboration can occur in a variety of scenarios. The ’time-space’ matrix [51] used in computer-supported cooperative work (CSCW) categorizes these scenarios according to whether users work in the same time and/ or space. In collaborative visualization, users can interact with visualization at the same time or not (synchronously or asynchronously). In relation to space, team members can be collocated or geographically distributed. While substantial previous research on CSCW and collaboration technology has focused on distributed and asynchronous collaboration, recently a lot of attention has been devoted to collocated collaboration.

Various collaborative scenarios fall under collocated situations and many others happen in a mixed presence [54]. In mixed presence, there is a mixture of collocated and distributed situations: some parts of the analysis happen in a distributed fashion but at least some parts happen in a collocated fashion. The collocated gathering provides analysts a better opportunity to share and discuss their findings. These collocated sessions are usually very time critical and team members need to make the best of their time [54]. For example, collocated collaboration is a frequent occurrence for intelligence analysis and business scenarios [21, 46, 55, 109]. Other well known examples of collocated scenarios include collaboration of medical practitioners to examine a patient's medical record, a team of geologists gathering around a large map to plan an upcoming expedition, a team of industrial designers to discus specific designs, a construction management team to coordinate the design of building systems, or a team of executives looking at charts showing the latest sales trends.

In addition to the importance and occurrence of collocated situations in real life, studying collocated scenarios helps to build a foundation to understand distributed and asynchronous collaboration. A lot of research in distributed collaboration aims to better understand how people interact during collocated situations in order to provide better support for distributed settings. For instance, Kraut et al. [59] discussed the effects of proximity on collaboration and how it could provide the foundation for a discussion of the actual and potential role of communications technology in distributed settings. Another example is O'hara et al.'s [78] description of a blended interaction space for small group distributed meetings that attempts to provide a sense of togetherness. There are many other instances such as Broughton et al.'s work [11]. Therefore, investigating collocated scenarios and understanding the challenges and possible leverage points to design collaborative tools seems critical not only for collocated scenarios, or a blend of collocated and distributed situations, but also as a fundamental understanding of teams' needs during collaboration to improve design of distributed tools. The latter is not the direct focus of this research, but I believe some of the findings can inform the collaborative VA community in general.

The benefits that collaboration offers in many disciplines, as well as the critical need to support collocated collaborative scenarios in real life, has inspired many researchers to focus on collocated collaborative VA situations (e.g. [31, 46, 48, 77, 94, 102,106]). However, providing tool support for information search, access, sharing and discussion during collocated collaborative VA is not a trivial task. These challenges make this research direction a very promising and interesting field of research.

2.2 Visual Analytics

Visual analytics (VA) is a discipline that studies the science of analytical reasoning. Cook and Thomas [18] defined visual analytics as “the science of analytical reasoning facilitated by interactive visual interfaces”. Visual analytics is an iterative process that involves information gathering, data preprocessing, knowledge representation, interaction and decision making. Keim et al [56] discussed the goal of VA as helping analysts “to gain insight in the problem at hand which is described by vast amounts of scientific, forensic or business data from heterogeneous sources”. To reach this goal, visual analytics combines the strengths of computer supported tools with those of humans.

Historically, visual analytics has evolved out of the fields of information and scientific visualization. According to Colin Ware [111], the term visualization can be defined as “a graphical representation of data or concepts”. Card et al. [12] defined visualization as “interactive visual representations of data to support human cognition”. Today, many software tools are employed to help analysts to organize information, create meaningful visualizations and explore the information space in order to extract potentially useful information [56]. However, according to Keim et al. [56] there are many challenges in this field. Therefore, we need to employ more intelligent means in the analysis process. While the ultimate goal of VA is allowing analysts to apply advanced computational capabilities to augment the discovery process, the transformation of data into meaningful visualizations is not a trivial task [56]. Very often, there are many different ways to represent the data and it is unclear which representation is the best one. In addition to challenges specific to designing VA tools, collaborative visual analytics involves another complicated factor: social processes [36]. Sensemaking is one of the most important human activities involved in VA. In the next section, I define sensemaking and discuss challenges to support the collaborative sensemaking process.

2.3 Collaborative Visual Analytics Process

While substantial research has been devoted to computer-supported cooperative work in general, collaborative visual analytics is still not fully explored. It is still not thoroughly clear how people collaborate to solve data analysis tasks, or how information visualization techniques and interaction methods need to change to better support collaborative work. Therefore, the first fundamental step is gaining a better understanding of how users analyze data to characterize the processes and activities involved (e.g. [29]). This kind of research can inform the VA community about possible challenges and design implications. The next step for this research is to test and evaluate proposed design implications in order to recommend more practical guidelines. Recently, some research has begun to address the first step: understanding the process by conducting either case studies in the field or controlled lab studies.

More relevant to my work are studies that consider analytic processes of groups [48, 50, 67, 81, 86] by using software supporting collaborative work [68, 81] or by using paper-based tasks [48, 86]. Findings of previous studies, regardless of whether the tasks were paper based or software based, resulted in similar lists of processes involved in collaborative data analysis. For instance, Mark and Kobsa [68] identified processes of parsing the question, mapping variables, finding or validating a visual representation, and validating the entire analytic process. Isenberg et al. [46] categorized analytic activities in a collaborative context. They derived eight primary visual analysis processes: browse, parse, discuss collaboration style, establish task strategy, clarify, select, operate and validate. Each process contains a number of activities. For instance, while browsing, participants scanned, flipped through and grouped visualizations to gain a better understanding of available information.

In a single user context, Gotz et al. [29] identified and categorized various visual analytic behaviours. Their four-tier hierarchy is comprised of tasks, subtasks, actions and events. They argue that the action layer carries information regarding users' analytic intention/s. With a narrower focus, Sarvghad et al. [91] compiled a list of the most probable history operations for collaborative settings (browse, search, filter, edit, delete and export).

My first observational study (Chapter 3), demonstrated that record-keeping is critical in collaborative analysis situations. Therefore, in the next section I discuss the state of the art in regards to record-keeping, its importance and available record-keeping support in the context of visual analytics.

2.4 Record-keeping, Externalization and Task History

Record-keeping enables users to offload their memory during the course of analysis and later review, revisit, and retrieve prior visualization states and analytical findings. In this thesis I refer to record-keeping as activities that involve recording any information related to the analysis task. Recorded information can include visualization snapshots, system states, notes, and annotations. Record-keeping enables users to create a history of previous actions and/ or steps taken as well as record the reasoning behind those steps/ actions. For instance, a user may record a hypothesis in a note and later use the note to explain why he or she came to that conclusion and what were the sources behind their reasoning. While record-keeping refers to the activities done by an analyst, I use the term task history to refer to the recorded materials themselves (e.g. recorded visualization snapshots, states, notes, etc).

Prior research often refer to note taking and/ or annotation as *externalization* and recorded visualization snapshots and system states as *history*. In contrast, my definition of record-keeping includes both externalization in the forms of notes/ annotations as well as visualization states. In this thesis I define externalization as the process of taking information from one's mind and representing it in an external artifact. My definition overlaps with previous definitions of externalization in terms of recording insights in the form of notes, annotations, and bookmarks. However, I want to emphasize that my definition of externalization not only includes externalizing insights but also offloading any information related to the course of analysis such as to-do lists or reminders for further analysis.

My findings in Chapter 3 highlight note taking as a pivotal activity during the course of analysis, emphasizing the importance of supporting note taking and integrating it with history tools. In the following subsections, I review previous literature on both externalization and history support. I will discuss the importance of supporting each of these forms of record-keeping and how they have been implemented in the context of VA for single users and collaborators. Later in Chapter 5 I discuss how creating links between different collaborators' findings and connecting common externalizations can exploit the record-keeping support.

2.4.1 The Importance of Externalization

According to Pirolli and Card [84] it is difficult for the human working memory to keep track of all findings. Synthesizing many different findings and relations between those findings increases the cognitive overload and hinders the reasoning process. Therefore, "External Cognition" (i.e. the notion of using external world for assisting thought and reasoning) enhances cognition [92]. Card et al. [12] discussed the results of an experiment in which multiplying two digit numbers using pen and paper (external aid) took significantly less time than doing the same multiplication mentally. They argue that the challenge is holding the partial results in memory rather than the multiplication itself. Pen and paper help with the former by enabling people to offload the partial results externally on paper. By extension, record-keeping, a form of external cognition, can assist visual data analysis and reasoning by providing analysts with means for recording, reviewing, validating and sharing their findings and insights.

Note taking has been the subject of investigation in many domains such as education, cognitive psychology, and visual analytics. It is used daily as an information processing tool for many different purposes [35]. From a psychological point of view, taking notes is a way of offloading cognitive processes and intellectual products such as insights, findings, and hypotheses. It helps to build a "stable external memory" that can be used at a later time [7]. Furthermore, note taking seems to assist complex tasks such as problem-solving and decision-making by reducing the load on working memory [7]. It has also been observed that taking notes keeps students engaged and improves the learning process [7, 10].

The importance of note taking and annotation in visual analytics has also been mentioned [37, 43, 49]. Externalization is known to support insight generation [37, 43, 49, 52, 61, 97, 98]. Heer et al. [37] stated that annotations and notes are important for supporting discussions around visualizations in distributed collaborative visual analytics. Kadivar et al. [52] mentioned that annotating visualizations can be effective in supporting exploratory visual data analysis. Therefore, both textual and graphical annotation on visualizations may be necessary. Notes help analysts to think through problems and to remember previous findings and cues [69]. Furthermore, notes help to create a link between the system and an analyst's cognitive processes [97]. Lipford et al. [61] stated that externalization improves recall at a later time, which helps analysts to "discuss their rationale and decision points more confidently and clearly". Kang and Stasko [55] suggested that supporting insight provenance[1] and sanity checking for intelligence analysis could save time during report-writing. Their field observations showed that analysts spent substantial time returning to original sources to find the supporting references and rationale behind their statements. Vogt et al. [109] and Pirolli and Card [84] similarly pointed out the need for recording findings, hypotheses and evidence.

Schematizing

Several studies signified the importance of schematizing results [15, 47, 54]; in other words, organizing results and other externalizations into a structured format. Kang et al. [54] discussed users' difficulties in making connections between entities and reported that structuring notes was critical to successful performance. For instance, several structured formats can be useful for intelligence analysis, including timelines, spreadsheets, lists, and networks [15,47]. Zhang [117] discussed the nature of external representations in cognition and mentions diagrams, graphs, and pictures as a few typical types of external representations. For meetings or tasks that require flexibility, such as brainstorming and collaborative design, freeform graphical input could be a better option to support flexibility [53, 74]. Other structures include casual loop diagrams, mind maps, diagrams, graphs, and pictures [53, 73, 74].

I expect that schematizing may be even more critical for collaborative work, since the structure may additionally help with communication. Moreover, I expect that integrating the schematic views of different collaborators will help to build common ground and make it easier to identify related findings; I examine this hypothesis in Chapter 5.

2.4.2 Implementation of Externalization Support

Externalization has been implemented in several research tools. Sense.us [39] allowed users to collaboratively analyze data, add annotations on top of a visualization, and write down their findings in notes attached to the visualization. Aruvi [97] similarly enabled users to take notes and link related notes to each other and the visualization. Harvest [98] had a note taking mechanism that automatically recommended notes most related to an analyst's current line of inquiry. Collaborative Annotation on Vi-

[1] Insight provenance is a process that involves understanding where insights and decisions came from; it is much easier with proper record-keeping practices and tools.

sualization (CAV) [22] allowed analysts to remotely collaborate and add annotations on top of their visualizations. Research by Shrinivasan and van Wijk [97] suggested that the proposed benefits of record-keeping can be better exploited if a history mechanism not only captures the analysis process and externalized insights but also has the means to create links and connect stored artifacts.

Note that the purpose of note taking can be different in distributed collaboration than in collocated work or single-user systems. For example, note taking in Sense.us is primarily designed to support online discussions around visualizations that cannot occur through face-to-face dialogue rather than to support recall of an individual's or group's findings.

2.4.3 The Importance of Task History Support

In visual analytics tools, record-keeping support is often implemented in the form of a history module that stores previous states of the system, including the visualizations that were generated. Many researchers have mentioned the advantages of history tools and their importance [36–38, 46, 72, 79, 80, 87]. According to Shneiderman [96], history tools can play an important part in the visualization process by enabling users to review, re-visit, and retrieve prior visualization states. As Heer et al. [37] noted, history tools can also be used to create a report or presentation after analysis is complete. They also suggested additionally recording past visualization states, and also suggested that history improves communication and dissemination of findings.

Isenberg and Carpendale [43] stated that while data analysis histories are necessary for individuals, they might be more important for collaborative tasks. It has been speculated that capturing individuals' analytic activities can help users maintain awareness of each other's work, particularly while shifting between loosely to tightly coupled collaboration styles. In addition, recorded material could be used at a later time to discuss or share interesting findings [43]. These speculations have not previously been empirically verified.

2.4.4 Implementation of History of states

Several single user VA tools provide general-purpose undo/redo operations, but this simplest form of history is inadequate for most complex VA tasks. Heer et al. [37] suggested additionally recording past visualization states. They also recommended that history improves communication and dissemination of findings. Single user VA tools have implemented variations of task history. For example, Heer et al. [37] integrated a graphical history module into Tableau software [101], that enabled users to visually browse, search, filter and reuse previously created visualizations. VisTrails [5] captured detailed information about scientific workflow, including data, visualizations, and the pipelines used to create the visualizations.

Collaborative use of recorded materials has been mainly investigated in the remote asynchronous context. Heer et al. [39] found that record-keeping facilitated view sharing, threaded discussions, and social navigation. Similarly, Many Eyes [108] is another web-based tool that enables bookmarking and sharing of views to support discussion.

In the collocated synchronous context, the closest research to this work is Cambiera [45], a tool that tracks each individual's history while they analyze a document corpus. Using colour-coding, past searches and documents are visually represented to increase users' awareness of each other's work. Somewhat less related is MemTable [41], a smart tabletop surface that captured and visually represented the table contents during meetings, including individual participation histories. However, MemTable was designed for more general types of meetings rather than visual data analysis.

Although previous work has postulated that history tools may be even more important for collaborative work [38, 67], little guidance is available to help build such tools effectively. Extending history mechanisms to represent activities of multiple collocated users is non-trivial owing to issues of awareness, disruption, organization, and so on.

I further emphasize that the vast majority of history/ provenance tools have focused on either single-user systems or distributed collaboration. To extend previous history tools to support collocated collaboration, we need to understand how these tools should be changed. For instance, one of the requirements of collocated collaboration is supporting transient collaboration style (see section 2.5.1). It has been speculated that capturing individuals' analytic activities in a history tool can help users maintain an awareness of each other's work, particularly while shifting from loosely to tightly coupled collaboration and vice versa [33]. In this regard, one of the challenges is how the balance between individual and collaborative work should be considered. Moreover, some of the current forms of record-keeping are inadequate for complex VA tasks. Therefore, we need to understand how history tools should be designed to better support VA tasks.

2.5 Design Considerations for Collocated Collaborative VA

There are many challenges and requirements unique to designing tools for collocated collaborative visual analytics [43]. Different bodies of work have addressed these challenges. As mentioned earlier, interfaces, visualizations, and interaction techniques should be designed to address the specific needs and requirements of collocated analysts. Scott et al. [93] suggested a list of guidelines that collocated collaborative technology must support. These guidelines include providing support for **(1) natural interpersonal interaction, (2) transitions between activities**, **(3) transitions between personal and group work**, **(4) transitions between tabletop collaboration and external work**, **(5) the use of physical objects**, **(6) accessing shared physical and digital objects**, **(7) flexible user arrangements**, and **(8) simultaneous user interactions**.

More recently, Isenberg [42] gathered a list of guidelines for collocated collaborative VA. Table 2.1 represent these guidelines. However, it should be noted that these design considerations were not available at the time of my first user study. At that time I had access to a less complete version of this list available in [43].

Through the process of my research the challenge of supporting record-keeping activities emerged. In addition, based on previous work providing support for transient collaboration style, communication and coordination and awareness are necessary to support collocated collaborative VA. Therefore in the following subsections, I provide more background on these challenges.

2.5.1 Transient Collaboration Style

Many collaborative tasks require changes in collaboration style, where people move back and forth between individual and group work [33]. According to Tang et al. [102], collaborators tend to frequently switch between loosely and closely coupled work styles when working over a tabletop. Another study [76], demonstrated that users preferred to work individually on some parts of a problem when the system was capable of supporting such individual activities. In an exploratory study, Isenberg et al. [46] identified eight types of collaboration styles that were used during a collaborative problem-solving task. This diversity highlights the necessity of supporting various

Consideration	**Aspects to consider**
	Collaborative Environment
Display Size	Socially appropriate work space size per person, establishment of private, group, and storage spaces
Display Configuration	Accommodation of group's current work practices, tasks, and goals
Input Type	Impact of input type on possible interactions
Resolution	Input and display resolution
	Supporting Social Interaction
Communication	Explicit data referencing across different representational and viewing contexts, e.g., annotation, implicit awareness cues of changes to the data across different representational and viewing contexts, support clarification, validation, data and strategy discussions with shared artifacts
Coordination	When using individual data views, location and rotation as a coordination and communication tool, sharing of visualizations and views, multiple synchronous interactions with shared representations, temporal flexibility of analysis processes, analysis processes within different collaboration styles, awareness of analysis activities and histories to encourage joint work
	Designing Information Visualization
Representation	Personal preferences, multiple representation types, awareness support, appropriateness of representation for work environment and social interaction, integrated reasoning and sensemaking support
Presentation	Arrangement of data items for group access, providing copies of the same data, accommodation of input methods, compensations for display resolution
View	Interpretability of data from multiple viewpoints and orientations
Interaction	Interactive response rates despite simultaneous interaction, collaborative interaction histories, conflict reduction arising from global changes to data or view, fluid interaction, temporal flexibility of analysis processes, support operation, selection, browsing, and parsing as parallel activities

Table 2.1: Design considerations for collocated collaborative data analysis environment (Isenberg [42])

collaboration styles. Moreover, it is also known that the nature of the task influences the collaboration style [36]. Cooperative tasks may encourage a lot of communication, while competitive tasks may encourage participants to work more individually. For example, in the second user study 3, participants in an observational study expressed a desire to work individually at the beginning of a competitive task and share their results towards the end.

Previous research has identified the need for flexible workspace organization for collocated work on interactive surfaces [94]. This is one known way to support changing collaboration styles. Creating copies of data representations enables analysts to work on a common data view in parallel. Scaling and rotation helps analysts to work on any side of a tabletop display and to orient visualizations in ways that accommodate concurrent analysis [43]. Lark [104] is one tool implemented for collocated collaborative analysis of data. A pipeline metaphor is used to organize the various views of the data and demonstrate their relationships, so that users can see how their own work relates to views created by others. Analysts can create several copies of views at different visualization stages. These copies can be used individually or can be shared by the group. While the pipeline metaphor has the benefit of explaining relationships between representations, it has the downside that the position of data views on the tabletop surface is constrained by the global pipeline. DTLens [25] is a tabletop system for exploring large maps and diagrams, which allows individuals to edit their own views of the data. Having separate visual representations prevents conflicts of interest that may arise when multiple people simultaneously interact with information.

2.5.2 Providing Awareness

Shared awareness (understanding who you are working with, what is being worked on, and how your actions will affect others) is critical to ensure efficient and effective team coordination and decision-making [60]. In collocated collaboration, people are able to gather implicit information about team members' activities from body language, alouds, and other consequential communications [38]. For these reasons, co-present situations may have fewer awareness challenges compared to distributed situations. Nonetheless, there are still issues to be considered. For instance, awareness becomes a challenge when group members are working in a loosely coupled fashion since conversation may be disruptive [43]. This is particularly true in complicated visual analytics tasks, where users can easily duplicate each other's work (e.g., by creating the same charts of a data set). Therefore, there should be channels for providing awareness with minimal interruption and cost.

Few visualization tools for collocated collaboration provide explicit mechanisms for awareness. With Lark [104], users can create several copies of data views at various points along a visually presented visualization pipeline. Changes at upstream locations (i.e. closer to the dataset on the pipeline) are propagated into all the downstream data views. Though it reveals the downstream changes, I believe, in line with the authors, that this approach works better for coordinating work rather than providing awareness of what others have done. Colour-coding has been used as a mechanism for providing awareness in collocated collaborative situations. With this approach, each collaborator is assigned an exclusive colour by the software that marks his or her analytical work with some visual cue of that colour. An example is Cambiera [45], a system designed for collaborative visual analysis of text documents. Each user's searches are marked with varying shades of one colour. This enables collaborators to recognize and track their own and each other's work.

2.5.3 Communication and Coordination

Various studies have found that the closeness of groups' collaboration styles directly affected outcomes [8, 47, 55, 109]. Isenberg et al. [47] and Vogt et al. [109] found that teams who collaborated more closely were more successful, and Bradel et al. [8] reported that members in more successful teams understood each other's work better. Groups should not always work in a closely coupled fashion, however. Kang and Stasko [55] found that in a long-term project, collaboration was loose during information collection but tight when synthesizing findings and writing a report. 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 to each investigator about their collaborators' activities and findings. Paul and Reddy [82] suggested providing support for both action awareness and activity awareness, showing actions that led to a particular activity.

Various techniques have been developed to help collaborators maintain awareness. For instance, Hugin [58] provided awareness support in a mixed-presence setting through a layer-based GUI design. Some of these, such as the tabs in CoSpaces (discussed in Chapter 4), 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 since they are often hidden from view. Nonetheless, the study on CoSpaces revealed that tabs were useful for exploring other people's work in a non-disruptive way.

In contrast, integrating everyone's information into one view could cause disruption to individual work as the view constantly updates. Brennan et al. [9] implemented a visualization of externalizations and explored different ways to merge collaborators' content. Similarly, CoMotion [16] 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. Another related system is Cambiera [45], which provided awareness cues about related searches conducted by a collaborator. Isenberg et al. [47] reported that these cues, which they termed *collaborative brushing and linking*, encouraged closer collaboration. I emphasize that Cambiera focused on information foraging; it did not support the sensemaking loop nor consider how linking could be applied to recorded notes and findings. My work extends the linking idea to externalizations. Note that I use the term 'linked common work' instead of 'collaborative brushing and linking' to generalize the linking idea to cases where there is no active 'brushing' action to initiate the link.

It is also important to support coordination to allow collaborators to coordinate their many tasks [68]. Mark and Kobsa [68] referred to supporting coordination as a challenge for designers and emphasize the need to choose appropriate visual representations to help teams during VA tasks.

2.5.4 Shared and Individual Workspaces

The importance of providing both individual and group workspace is well known [33, 45, 47]. For example, Wallace et al. [110] demonstrated that an investigative task had better outcomes when the group could share items through a tabletop display than when they had only tablets.

This need for shared and individual workspaces applies equally well to externalization spaces. Even though collocated teams often assign one note taker [86, 109], my first study showed that individuals still occasionally need to record private notes (Chapter 3). This suggests that it is important to provide both shared and individual note taking spaces. This study also revealed that taking notes on paper reduced the note taker's awareness of other activities, suggesting that note taking tools should be integrated with the investigative application. However, a separated digital view can suffer from the same problem. For instance, in Bradel et al.'s [8] study, only one user could take notes in the shared space; others had to work on separate views and this reduced awareness.

Kang and Stasko [55] recommended "promoting individual workspaces as well as the ability to share sources and data, view and comment on others' work, and merge individual work together." Brennan et al. [9] similarly suggested providing individual perspectives on a shared common ground space. McGrath et al. [71] introduced Branch-Explore-Merge, a useful structure for transitioning between individual and group workspaces. I note that their concept was applied and evaluated for direct views of data, not within a collaborative thinking space. With CLIP, I address the need for individual workspace, as well as providing each user with flexible control over how much of the collaborators' information is shown in their view.

In the third phase of this research, as described in Chapter 5, I focus on providing support for collaborative sensemaking. To better understand the collaborative sensemaking challenges and requirements, I address an overall review of this topic in the next section.

2.6 Sensemaking

Sensemaking is a creative and retrospective process of deriving meaning from information to help in making decisions [113]. Similar to any creative process, it is not structured and requires interaction with the data to discover interesting facts [27]. Sensemaking is a challenging process that requires iterative detecting and adjusting of patterns, where an analyst creates hypotheses and validates them through various interactions. Collaborative sensemaking introduces a new set of challenges. In particular, effective collaboration requires effective coordination and communication. Coordination is a multifaceted process that includes division of labor, sharing of information, and discussion of hypotheses to help the team come to a consensus [2]. This process is difficult for most collaborators and there are many ways to fail. In addition, communication is an integral part of collaboration. Therefore, providing support for coordination and communication is essential for effective collaborative sensemaking.

To provide better support for collaborative problem solving, it is important to understand how analysts tackle complex tasks such as intelligence analysis or business problems. In this section, I review the sensemaking models which describe how an analyst searches for and organizes data, and creates new knowledge and theories, in order to solve a problem.

Russell et al [89] presented a sensemaking model that contains two top level iterative processes. The first is searching for a representation where a suitable frame for encompassing information is considered. The second is encoding the information into the representation and adjusting the representation to accommodate items that do not fit. Card et al. [12], suggested a high-level sensemaking model. Based on this model, the iterative process of sensemaking consists of five major activities: data foraging, searching for a schema, instantiating a schema, problem solving, authoring and acting.

In this thesis I refer to a more recent model by Pirolli and Card [84], which was developed in the context of intelligence analysis. This model demonstrates that analysts make connections through two iterative loops: information foraging and sensemaking (Figure 2.1). The information foraging loop involves searching, reading, filtering and extracting information, whereas the sensemaking loop involves iterative development of a mental model that leads to formation and reevaluation of hypotheses, and publishing the results. Pirolli and Card [84] described the sensemaking loop as a high cognitive load process that is affected by many cognitive biases: First is the limitation of the human's working memory to remember hypotheses, evidence and evidentiary relations. Therefore, information visualization techniques can play an important role to expand analysts' working memory by offering external memory. In addition, human perception is biased towards interpreting information in ways that support existing expectations. Thus, during problem solving and decision making tasks, people may fail to generate new hypotheses. This is especially true under time pressure and data overload. Another way to improve this process is employing methods for generating, managing and evaluating hypotheses. Confirmation bias is when people fail to confirm or disconfirm hypotheses. Pirolli and Card [84] suggested that distributing more attention to highly diagnostic evidence will help people to search for disconfirming relations. Therefore, representing a set of generated hypotheses and most importantly representing evidence can help analysts to better connect the dots and make a decision.

While support for collaborative sensemaking has been broadly investigated due

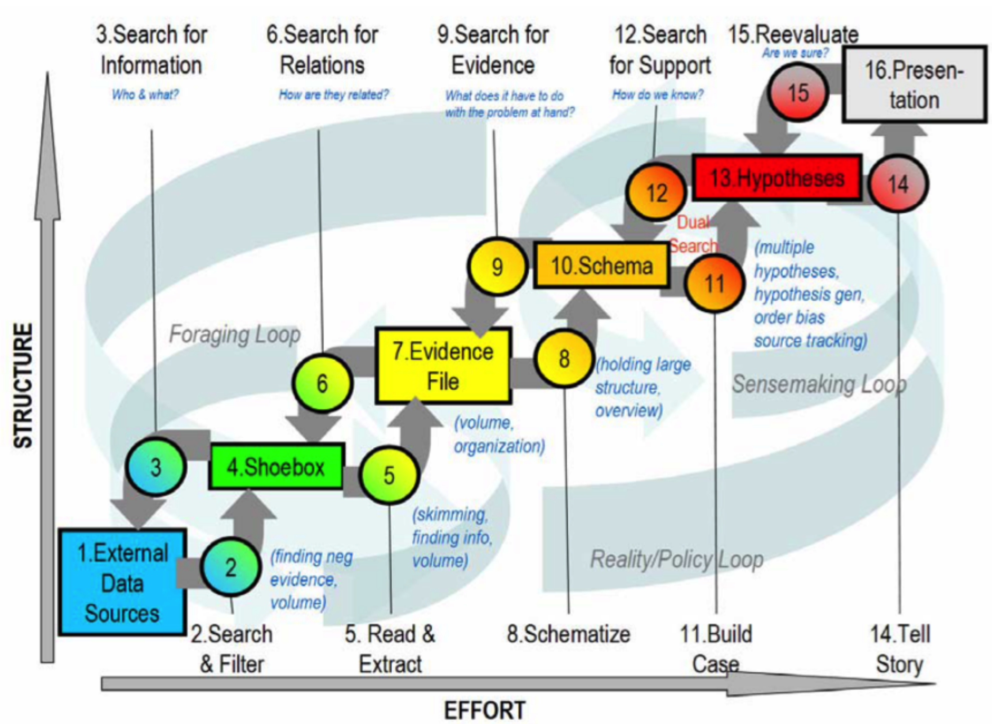

Figure 2.1: Pirolli and Card sensemaking model [84].

to its growing importance in solving complex tasks, developing appropriate tools to support the sensemaking process is still a challenge. Support for collaborative sensemaking can be more challenging than for a single user. This is mainly because, in addition to handling various set of complex tasks involved in sensemaking, collaborators need to coordinate their actions together. In collaboration, the social aspect of sensemaking adds more complexity to the design of sensemaking support tools. Willet et al. [115] suggested that collaborative sensemaking tools should support group exploration and evidence gathering tasks by helping users build on each others' findings and collaboratively organize and synthesize them. While this suggests supporting the sensemaking loop, there is still a lack of support for the sensemaking loop in collocated collaboration. Well-known visual analytic tools such as Jigsaw [100] focused on the information foraging loop and offer limited support for the sensemaking loop. Tools with more focus on the sensemaking loop have been either designed for individual use or distributed collaboration.

2.6.1 Studies of Collaborative Thinking Spaces

Most similar to my work in Chapter 5 are studies of collaborative visual analytics work involving thinking spaces. Most important are studies that explore how to integrate or provide access to different users' views and what level of integration is appropriate for different kinds of shared information.

Chen et al. [13] built a tool that enabled asynchronous collaborators to record and share insights, and demonstrated that people could learn from others' past insights. Similarly, Willett et al. [115] demonstrated that asynchronous collaborators benefited from the ability to classify their text comments using tags and from the ability to link comments and views; however, in their system, users had to create links manually. Neither of these studies examined the value of automatically linking common work, nor did they look at synchronous collaboration.

For synchronous collaboration, most evidence suggests that highly integrated views should be a good approach. Balakrishnan et al. found that a shared view of a network diagram supported better performance at a collaborative intelligence task than separated views [3]. I note, however, that their network diagram was a simple social network rather than a dynamic thinking space for recording and organizing externalizations. I am motivated by their later study [4], where they suggested that it might help to additionally visualize partners' activities and externalizations. A study [8] using Jigsaw's tablet view did allow analysts to record their findings in the form of notes, graphs, and timelines. Bradel et al. [8] found that Jigsaw's tablet view was inadequate for collocated collaboration, because participants wanted to use it as a shared note space, but it accepted only one person's input at a time. Their study did not examine separate linked views of the thinking space, but suggested that it may be a good idea.

2.7 Summary

While there is an apparent need for tools that support collaborative sensemaking and reasoning, our knowledge of how analysts actually use these tools and possible challenges is limited. In this thesis, I aim to increase our understanding of how analysts use such support tools and how we can better support this complex process.

To this end, I conducted a series of user studies that resulted in richer understanding of the collocated collaborative process and several guidelines for supporting record-keeping and note taking (Chapter 3). In Chapter 4 I discuss the effect of integrating a record-keeping module with a VA tool. Later in Chapter 5, I demonstrate that automatically linking collaborators' findings has value to increase awareness during the collaborative sensemaking process. Linking collaborators' findings also resulted in more effective coordination of work and increased discussions around findings and hypotheses.

Chapter 3

Phase 1: Understanding Collocated Collaborative Visual Analytics Processes and the Role of Record-keeping

This chapter addresses the first research question: **How can collaborative record-keeping activities be characterized in the process of visual analytics for a small team of collocated collaborators?**

In Chapter 2 I discussed the complexity of the collocated collaborative analysis process and the critical need to better understand the process, the challenges involved and the requirements to support these challenges. This understanding is crucial in order to propose practical guidelines. To this end, I conducted an observational laboratory study to examine the process of collocated visual analytics in a business context. [1] The study was exploratory in nature rather than designed to test a specific hypothesis. My goal was to better understand collaborative activities and challenges that might suggest improvements for collaborative analytics tools.

Better support for record-keeping activities emerged as a critical need for collocated collaborative VA during this study. This finding formed the direction of this thesis and prompted me to analyze note taking (paper based, in this study) and other record-keeping activities in depth in order to establish design requirements. In this

[1] This study was published at the VAST 2010 conference [62] with an extended version in the Information Visualization Journal [63].

chapter, I propose a characterization of VA activities highlighting the importance of record-keeping during this process. In addition, I present a categorization of notes based on their content, scope, and usage, and discuss how record-keeping fits into the visual analytics process.

In Chapter 2 I discussed related work on collaborative visual analytics process. For instance, Mark and Kobsa [68] identified processes of parsing the question, mapping variables, finding or validating a visual representation, and validating the entire analytic process. Isenberg et al. [48] identified processes of browsing, parsing, discussing collaboration style, establishing task strategy, clarification, selecting, operating, and validating. These lists of activities bear strong resemblance to my own characterization. For instance, the first phase that I defined, "problem definition," has been identified in previous work as "parsing" [48,68] or "problem interpretation" [81].

In contrast to previous work, my framework captures the whole process of visual analytics (as opposed to, say, simply the visualization or analysis phase) and breaks each phase down into lower level activities. I believe this two-level structure provides a useful way to think about the analysis process. I also highlight record keeping as a critical activity during all phases. Although record keeping has been previously mentioned as a relevant action [12, 39, 96], its importance may have been underrepresented in previous frameworks describing collaborative analytics processes.

The main contributions of this chapter are recognizing the critical role of record-keeping in the overall process of collaborative visual data analysis and discussion and design suggestions around how record-keeping could be integrated into collaborative visual analytics tools (described in Chapter 1 as C1, C2 and C3).

3.1 Introduction

Analysts rely heavily on insights discovered in the course of data analysis. These insights are shared with others and used to assist with higher-level tasks such as decision-making and problem-solving. To help users remember, share, and make sense of their insights, Lipford et al. [61] argue that there is a critical need to support externalization through mechanisms such as taking notes, and annotating views. However, support for these activities in visualization tools is currently limited. In this chapter, I discuss the results of an observational user study that emphasized the significance of record-keeping activities during collaborative visual analytics on interactive surfaces. My focus is on collocated work by small groups of known collaborators, as illustrated in Figure 3.1.

Figure 3.1: Examples of note taking activities during the observational study. Sometimes note-takers are disconnected from group activities.

Visual analytics tools are becoming prevalent in a variety of domains, including business. These tools assist users to examine complex datasets and interactively explore relationships and trends [38, 61]. As a result, business intelligence tools have been widely adopted. These tools are typically designed for single users working on desktop machines, whereas business tasks often require users to work collaboratively, particularly when each user has unique expertise or responsibilities.

There is growing interest in developing visual analytics tools for large touch-sensitive wall and tabletop displays, a potential solution to the collaboration problem. Such display technology is known to facilitate collaborative work by allowing users to interact and explore a dataset simultaneously [38, 43]. Particularly the physical affordances of the tabletop display offer the potential to enhance collaboration by encouraging group members to switch roles, explore more ideas, and follow each other's actions more closely [87]. However, collocated collaborative visualization is a relatively new area of research and still not fully explored. There are many special challenges and requirements for the design of collocated collaborative visual analytics tools and interaction techniques [43]. Different bodies of work have addressed some of these requirements, but only a few collocated collaborative visualization tools have been developed (e.g. [24, 25, 43, 45, 104, 107]). Other systems implemented for collaborative visualization, such as [37, 98, 108], have focused on distributed work, which has different requirements. None of these tools so far has focused on record-keeping activities during collaboration.

3.2 Observational Study

This laboratory study was exploratory in nature rather than designed to test a specific hypothesis. Groups of users answered focused business questions and participated in a competitive business scenario. Their work was supported by visualizations of sales data on large interactive wall and tabletop displays. In Chapter 1 section 1.3 I described the choice of method for this study including why I chose a lab study over a field study. Specifically, my goal was to gain a better understanding of collaborative activities and challenges that might suggest improvements for collaborative VA tools. However, the use of large interactive surfaces for this study (which were rarely used in practice at the time) made it almost impossible to conduct the study in the field. Details about this study including consent form, introduction to the task and system, tasks, and follow up interview questions can be found in Appendix A.

3.2.1 Participants

Twenty-seven student participants took part in this study, divided into nine groups of three. To simulate common work situations, all the group members were required to know each other. To mitigate the possible impacts of using students, I mainly selected participants (seven out of nine groups) who were familiar with the business domain (advanced BCom or MBA students). Participants of the other two groups were computer science graduates. All participants had experience with some kind of data analysis software such as Microsoft Excel.

3.2.2 Dataset and Task

Each group completed two tasks, both using a sample e-fashion dataset from Explorer [90]. The dataset contained information about sales of garments in eight states of the USA for three consecutive years. It consisted of nine columns and 3273 rows of data.

Task 1 consisted of six focused questions designed to help users learn important features of the visualization software. An example question was “How does the 2003 margin compare to previous years?” Task 1 was intended primarily to help users become familiar with the system. In Task 2, participants were asked to assume the roles of three managers (representing three different states) and together determine a marketing budget for the next year. They were told that rationale for the budget should be based on information within the dataset.

I received advice from business professionals and faculty members in designing the tasks. Because participants did not own the data and were not familiar with it, I was concerned about their engagement. Hence, I decided to first familiarize them with the data and the interface by using focused questions in Task 1. I then designed Task 2 to involve competition among the members of a group, which I hoped would engage the participants in the analysis process. Engagement was high, but the competitive nature of Task 2 did have an impact on the analysis process, as described later. Appendix A

3.2.3 Apparatus and Software

Identical rear-projected Smart DViT (digital vision touch) screens were used, one in a wall configuration and the other in a tabletop. Both had a size of 61.2 inches x 34.4 inches (70 inch diagonal) and had four high-definition projectors to create a total resolution of 3840 x 2160.

During the pilot studies I noticed that participants took notes on the margin or blank back of the 8.5 inches x 11 inches-sized instruction sheet; therefore, I decided to provide them with pens. I put pens on a table nearby and informed participants that they were available if required. However, at this point I was not focusing on note taking as the main interest and did not realize how important it would turn out to be.

Participants created and examined charts of data using"Explorer" [90] (Figure 3.2). Explorer allows users to interactively browse data, including selecting variables, filtering, and creating different types of charts. The tool was maximized to fill the screen. The software supported only a single input but each user had their own stylus so that they did not have to share a stylus to interact with the system. Note that Explorer was developed as a single-user application. I therefore expected some problems when using it collaboratively. I hoped that observing these problems would suggest changes that would better support group work. Using existing software enabled us to conduct preliminary requirements analysis without first designing a collaborative system.

3.2.4 Procedure

Prior to the user studies I run three pilot studies. This was mainly to ensure that 1) the analysis tasks were clear and understandable, 2) analysis tool was appropriate for

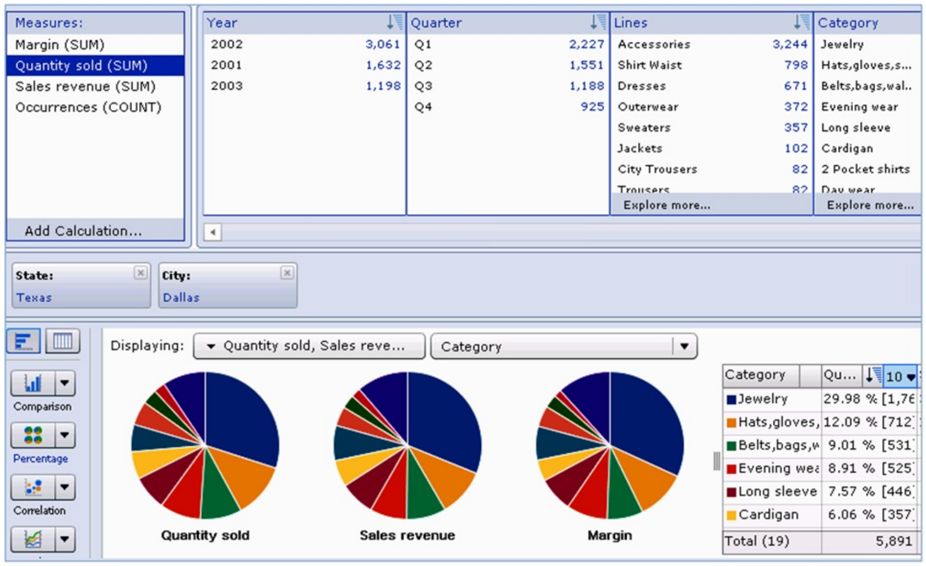

Figure 3.2: Partial screenshot of Explorer, depicting a comparison chart that visualizes margin, quantity sold, and sales revenue over category, filtered according to state (Texas) and city (Dallas).

the task and discover possible bugs, 3) the introduction, follow up interview questions and other instructions were adequate and unambiguous. During these pilots, I also determined ideal times for the different steps.

I began each study session with a 10-15 min introduction to the system (Explorer). Participants then spent approximately 30 min on Task 1 and 40 min on Task 2. I offered an optional 5 min break between the two tasks. After Task 2, participants spent approximately 10 min summarizing and writing down their results. I asked participants to create a report of their results at the end of Task 2 to justify their decisions. Following the computer-based tasks, I conducted an open-ended interview with all participants simultaneously.

Four groups used a tabletop display, four used a wall display, and the ninth group used both. This gave us an opportunity to obtain users' feedback on a variety of display configurations. Participants were allowed to arrange themselves freely around the displays, but generally had to stand to interact with them. Chairs were available near the tabletop and two large sofas were available near the wall display where they could sit if desired. A standing-height table was placed near each display with a mouse and keyboard that could be used if needed as well as the styli for direct interaction with the display.

3.2.5 Data Capture and Analysis

Data was gathered in the form of recorded videos, interviews, participants' paper notes, screen logs, and observations made by a live observer. In total, $\sim$ 630 min of video and screen logs were captured ($\sim$ 70 min per session) plus $\sim$ 20 min of interviews per session.

My colleague[2] and I coded and analyzed the data together using a structured qualitative data analysis process. We re-coded and re-analyzed data iteratively to refine the characterizations of collaborative processes, activities, and notes. In the first pass of analysis, we created form "I" to record chart types, values mapped to axes, filters, and time stamps from captured screen logs. Videos were used to fill in form "II", which contained positions of participants around displays. In form "III", we recorded the activities and events that each participant was engaged in during the analysis session. The main observation in the first pass of analysis was the importance and prevalence of record-keeping. Therefore, in the second pass we recorded the time stamp, purpose, and analysis phase in which participants took notes or saved charts. We refined forms "I" and "III" to capture this information. In the third and fourth passes we extensively analyzed and grouped users' notes based on their usage, scope, and content. Moreover, the roles of each group member (i.e. note-taker, software controller or observer) and role changes were recorded and added to form "III" in the fourth pass. Cumulatively, we spent nearly 2 months analyzing gathered information. Interview material was used to support and explain observations from the recorded material.

Note that because the study was not initially designed to focus on note taking, we were not able to fully capture the state of the notes throughout the study session, but only the state at the end. Therefore, we could not analyze how the notes varied between different phases. A different study design (e.g. using digital paper) could capture the note progression, and thus more completely enable such an analysis.

3.3 Findings

As the study was exploratory, I did not have any particular hypotheses. However, I had anticipated problems such as incorrect software orientation and inadequate

[2]This research was result of collaboration with my colleague, Ali Sarvghad. Thus any referral to "we" in this section acknowledges his work.

awareness of other users' work. I predicted these challenges to be the most important barriers and planned to concentrate on them, but based on my observations I found note taking a more interesting and yet less discussed obstacle to investigate.

3.3.1 Participants' Collaboration and Use of Software

I observed that group members were actively engaged in the analysis process. Their analysis activities mainly consisted of mapping and filtering data for new charts and having discussions about them. At any given time, only one of the group members was controlling the software, but they all participated in the cognitive process of analyzing the data. At times, users took turns to obtain information that was needed individually. Analyzing users' positions confirmed Tang et al.'s [102] results. Participants positioned themselves close to application controls and areas containing information such as a legend. At any time, the user who was closest to the widgets controlled the application; participants changed positions to allow one another to interact with the system. One of the impacts of the software not being designed for large screens was non-equal interaction. For instance, usually one person had to stand at a far corner to see a chart's legend and read it to the group. This likely impacted the groups' work style (i.e. led to closely coupled work with little parallelism), though I cannot be sure of the magnitude and significance of these effects. As the software layout dictated positioning, I did not examine position data in further depth.

Typically, one member of the group assumed the role of note-taker. Unlike role divisions observed in prior research [98], the assignment of roles was usually not discussed explicitly. Instead, typically one person would start taking notes, implicitly designating himself as the note-taker. Usually the act of one person writing something down after a group interaction or discussion would signal that a person was writing down something for the group. At times others would reassure themselves of recording an important item by asking the writer if he or she was putting it down. When I asked some groups about how they came up with their task division, they said it was based on their knowledge of each other's abilities. Role assignments typically remained the same throughout the work session, but in one instance the note-taker changed part-way through. In some sessions where one person was in charge of note taking, other members also took notes for themselves separately even though it meant that they had to stop working. For instance, participants of Group 7 completely stopped working seven times in Task 2 because they were all taking personal notes.

This clearly demonstrates that participants need to take notes both individually and separately from the group.

I did not observe any significant differences in the collaborative behaviour of computer science and business students, yet because of the small number of computer science groups (two groups), this observation is not generalizable.

3.3.2 Phases and Activities

I noticed a similar analytics process among all of the groups that I observed. I characterized groups' actions at two levels: high-level phases and low-level activities, as shown in Figure 3.3. As this characterization is grounded by the particular data that I collected, I cannot generalize the phases and activities to other data analysis situations. However, similarity to other frameworks [12,46,68,81] suggests that many aspects of this process probably occur outside the context of this study.

I identified four high-level phases: **problem definition**, **visualization**, **analysis**, and **dissemination**. Figure 3.3 shows that there are activities common to all phases such as record-keeping and validation, and activities unique to each phase. I explain each phase and their exclusive activities below. Common activities are explained separately.

My findings confirmed what Isenberg et al. [48] stated about the non-linear temporal order of activities. In this study, I observed that visualization, analysis, and dissemination occurred in a variety of orders, and the phases were re-visited multiple times within a session. This inconsistency in order confirms that flexibility is a critical design consideration.

Visualization and analysis phases were strongly interrelated, and participants moved back and forth between these two phases quite often. In the dissemination phase, participants returned to previous phases (e.g. to create a chart to include in their report) but with a lower frequency.

Phase I: Problem Definition

Users always started by building a common understanding. For example, they parsed the written description of the problem to make sure that they all understood what they were about to investigate, or they posed a new question to be answered. Having a consensus on what problem they were solving was the first step in working collaboratively towards a solution.

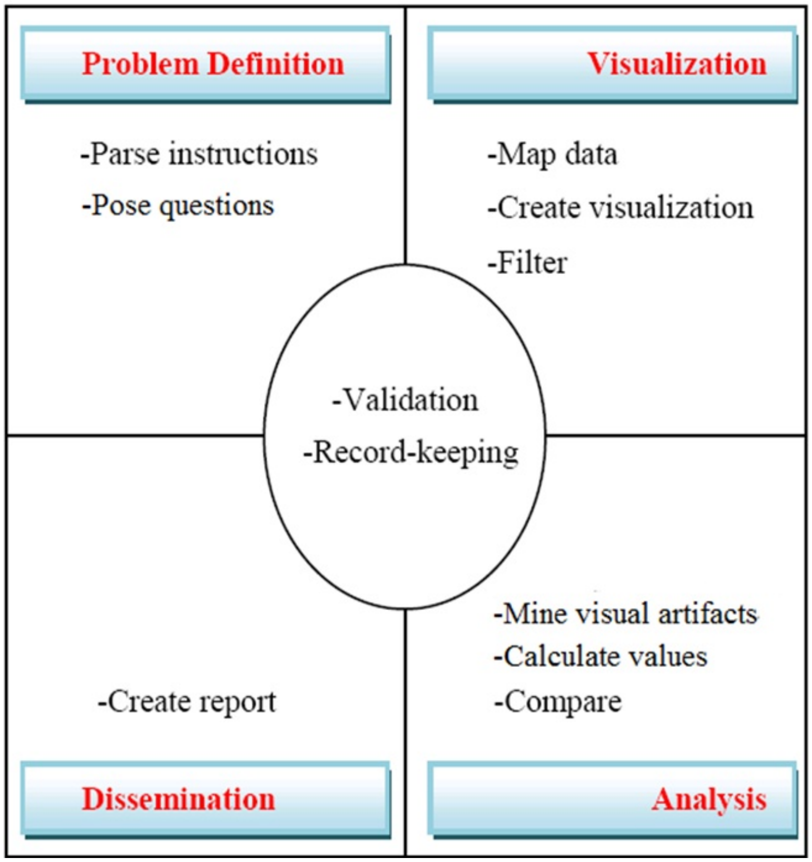

Figure 3.3: Activities I observed within the problem definition, visualization, analysis, and dissemination phases of collaborative visual analytics. Activities in the middle circle are common in all four phases.

Phase II: Visualization

The visualization phase comprises a number of essential activities, specifically:

- mapping variables;
- filtering underlying data;
- creating a visual artifact

Mapping variables and filtering data are steps in which users define variables, determine how variables will be visually represented, and extract a subset of data relevant to the problem. The final products of this phase are visual artifacts, which in this case were different types of charts. For instance, in order to reveal the trend of sales revenue in 2003, participants discovered that they needed to examine the values of sales revenue for all the quarters of 2003. Then they mapped “measure” to sales revenue and “dimension” to quarter. Next they chose a “correlation” chart to see the trend of sales revenue for 2001 to 2003. Finally, they applied filtering so that only 2003 data were shown.

Phase III: Analysis

Analysis is a complex phase that included activities such as:

- examining visual artifacts;
- making comparisons by referring to historical information such as notes or saved visual artifacts;
- calculating derived values through mathematic or statistical operations;
- gathering information from external resources.

The most common activity in this phase was examining charts. Participants worked together to extract information from the chart. In order to achieve this, they made comparisons, performed calculations, or searched for information through external resources. For instance, based on a chart depicting sales revenue for the four quarters of 2003 for California, one group decided that they needed to create similar charts for 2002 and 2001; then, they calculated sales revenue averages for all three years and compared values. They repeated the same activities for Texas and New York to make a decision on budget allocation.

Gathering extra information refers to the activity of accumulating necessary information to help address the problem. Extra information here means information that cannot be found in the dataset. The source of extra information could be a participant’s prior knowledge, the knowledge from Internet, etc. Figure 3.4 shows how a participant took note of values and extra information that could help analysis. “Texas: Dallas + Austin + Houston → huge market, oil = rich people” in the note is an example of extra information recorded by a participant, indicating that they thought there was a large market in Texas owing to the number of rich people who live there. Another example of extra knowledge was participants’ knowledge of the climate in different regions. They related higher T-shirt sales in Texas to the warm weather.

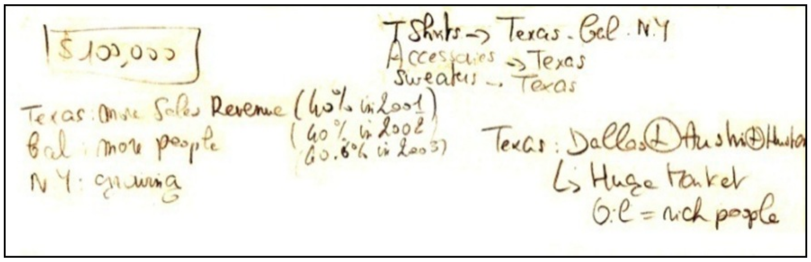

Figure 3.4: Note taken by a participant consisting of values such as $100,000, calculated numbers such as 60%, and externally acquired information such as “Texas: Dallas + Austin + Houston → Huge market.”

Usually, the product of this phase was a decision, an answer to a problem, or a hypothesis. In this phase, collaborators often carried out substantial discussion and negotiation to reach a consensus.

Phase IV: Dissemination

In the dissemination phase, participants used products of the analysis phase to generate a semi-formal report of their results. Reporting and presenting are very common activities in business, such as presenting results to the chief executive officer. I observed that while participants were preparing a report, they went back to previous phases. This usually happened when they were validating report material or providing extra content for the report, such as a chart or value.

Common Activities

Validation Validation activities occurred throughout the entire process and were concerned with ensuring the correctness of results and a common understanding. In the problem definition phase, participants verified a common understanding of the problem by posing questions. In the visualization phase, they verified the correctness of a chart by double-checking filtering and mapping of variables. In the analysis phase, they validated the acceptability of a budget allocation by re-examining charts. In the dissemination phase, participants checked the content of their final report to ensure that they were presenting the correct material.

Record-keeping In visual analytics, record-keeping ranges from visualization states to notes and annotations. In this study, this information took the form of charts saved by participants or notes that were written down on the instruction sheet. Participants took notes in the first two phases to define the strategy, saved values, and charts during the analysis and dissemination phases, and referred to their notes and saved charts to facilitate analysis and report writing in the dissemination phase.

3.3.3 Record-keeping Strategies

Table 3.1 reports quantitative data derived from my analysis of note taking and chart-saving for each group. This includes whether or not a group had a designated note-taker and the total number of people who took notes (including the designated note-taker). Table 3.1 shows that the dominant record-keeping activity in Task 1 was note taking. This can be attributed to several factors. First, the questions were more limited and could be answered with only one visualization rather than a series. Second, Task 1 did not require users to create a report. Another notable difference between tasks was that when there was a designated note-taker in Task 1, others did not take notes at all. In contrast, in Task 2, people often took notes in addition to the group notes taken by the designated note-taker. Similarly, in Task 2, often all group members took notes (six groups out of nine), whereas this happened in Task 1 for only two groups. A likely explanation for these differences is the competitive nature of Task 2 as compared with the cooperative nature of Task 1.

	Group	1	2	3	4	5	6	7	8	9
	Display used	Table				Wall				Both
Task 1	Number of note-taking actions	5	2	4	3	0	13	18	0	5
	Number of charts saved	0	0	0	2	2	0	4	0	2
	Designated note-taker	N	N	N	N	N	Y	N	N	Y
	Number of note-takers	2	1	1	3	0	1	3	0	1
	Number of times all took notes simultaneously	0	0	0	1	0	0	8	0	0
Task 2	Number of note-taking actions	8	4	7	7	20	8	11	9	6
	Number of charts saved	4	22	8	12	3	8	7	2	2
	Designated note-taker	N	N	N	N	Y	N	N	Y	Y
	Number of note-takers	3	1	3	3	3	3	3	1	2
	Number of group notes	0	0	0	1	1	0	0	1	1

Table 3.1: Number of note taking and chart-saving actions by each group. Shaded groups relied heavily on saved charts for analysis (N: no; Y: yes).

My observations of record-keeping strategies for Task 2 showed that groups could be divided according to three main approaches: five groups relied heavily on taking notes and saved only a few charts; two groups saved many charts and took few notes; and two other groups recorded nearly equal numbers of charts and notes. I focus only on the two extreme approaches. I believe that note taking and chart-saving approaches can be considered and studied as two different strategies for record-keeping for further analytic use. Note that the prevalence of the note taking strategy over chart-saving may be an artifact of my experiment as the process of saving charts was rather cumbersome. Because Explorer was not built with record-keeping as a focus, chart-saving was inconvenient and required users to select menu items and choose between various options. Charts could be saved only as non-interactive images.

Two groups selected chart-saving as their main strategy for keeping important information. One of the groups saved all the charts that they created during Task 2 and at the end they created a separate word-processing document in which they put all the charts side by side or further analysis. One of the participants of this group said, "I wish we could have all the charts on-screen to see them side by side," which implies that the tool used for analysis should have provided them with this functionality. The other group just saved a number of charts that they thought were more important. At the end they opened charts one by one for further analysis. Note that other groups saved charts as well, but less often and mainly for use in their reports.

Participants re-used the saved charts mainly for two purposes. One was for creating a report at the end of the analysis session (seven groups) and the other was for further analysis of data towards end of the analysis session, after creating several charts (two groups, shaded gray in Table 3.1). Groups who saved charts for the second purpose saved a larger number of charts. I cannot exactly pinpoint the criteria that different groups used to agree on the importance of a chart. Future work is needed to determine if there are any factors that can predict whether a chart is important enough to be saved.

The chart-saving strategy suggests that tools should enable users to save important artifacts and re-use them (e.g. as a history). History items may also reduce the number of notes that need to be taken because many findings are already recorded in the data representations. I noticed that the overall amount of note content taken by participants in Group 2 was less than that for the other groups. Based on my analysis of their activities, I attribute this to the fact that Group 2 saved many charts as image files. Note that these users still took some personal notes, so the ability to save charts does not eradicate the need to take notes. It seems that even the most sophisticated history mechanism is incomplete if it does not provide users with the ability to take notes.

3.3.4 Characterization of Note Taking Activities

Notes' content Based on my analysis of the notes taken by participants, I broke down a note's high-level content into *findings* and *cues*. Findings were recorded results of mathematic or statistical operations (e.g. 27% higher sales in New York, California's revenue is $60,000), observations (e.g. menswear sales are higher than womenswear in a graph), and decisions or outcomes of the analysis process (e.g. allocating more budget to Texas). A cue is anything noted by the user that is not directly extracted from a visual representation. For instance, users in this study wrote their interpretations of the questions in a concise form for themselves, or they drew circles around keywords in the instructions. Cues could be in the form of to-do lists or questions to be asked/ answered later on. For instance, one participant who had assumed the role of California's manager noted "T-shirts" as reminder to look into California's sales of T-shirts later on. Findings were also sometimes stored as saved charts rather than written notes. I noticed that in Task 1 (in which most users were not saving charts), the amount of note taking was higher than in Task 2 (in which users were saving charts). With respect to the visual analytics process described earlier, I observed that findings were mostly recorded during the analysis phase whereas cues were mostly taken during visualization.

At a lower level, notes typically contained one or more of the following elements: numbers (e.g. data values), drawings (e.g. flags, charts), text (e.g. questions, hypotheses, reminders), and symbols (e.g. %, $). In addition to ordinary use of symbols (such as $ for monetary values), participants used symbols to accelerate the note taking process and thereby decrease distraction from the main task. For instance, they used $\uparrow$ to indicate the increase of a value such as revenue.

Notes' Scope Based on the way that notes were shared, I divide them into group and personal notes. I consider a note as personal when it is taken for individual use, and as a group note when the writer intends to share it with the group. Personal notes were not necessarily private – in some cases, they were shared. For instance, during Task 2, participants shared personal notes to justify the amount of the budget they were demanding for their state.

I noticed that the nature of the problem influenced the scope of notes. During

Task 1 (which had a coopera- tive nature), usually one participant took notes for the group. In contrast, in Task 2 (which had a competitive nature), participants tended to take notes individually during the analysis phase (in addition to any group notes) and then referred to these personal notes during budget negotiations. However, based on my analysis of the group notes, it appears that designated group note-takers did not take additional personal notes, perhaps because they were already too busy.

Figure 3.5 shows a group note. It contains calculated values and has been nicely formatted to allow users to compare certain values of interest across the three states. This was used as a summary to help decide budget allocation. The tabular data made the analysis task easier by saving important information and enabling comparison; it seemed more convenient and efficient to record this information than to re-visit previously created charts. The same person who was in charge of note taking also created the final report.

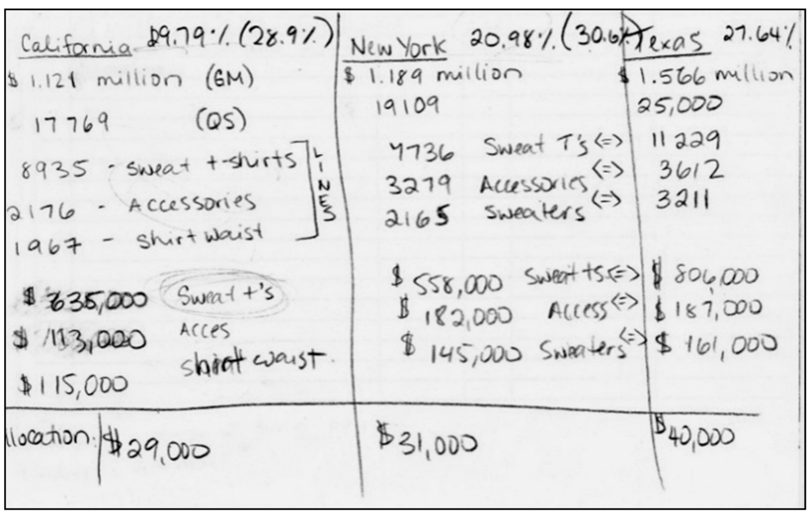

Figure 3.5: A note that was taken for group use. It is nicely structured and comprises information for all three rivaling participants.

In contrast, Figure 3.6 shows personal notes taken by three participants of a group. It can be clearly seen that the notes are less structured and every participant just recorded what they found important to themselves. For example, a participant wrote down sales revenues for 2001 to 2003 for his state of interest, which was California (Figure 3.6(a)). Similarly, other participants also noted interesting findings for their own states (Figure 3.6(b and c)).

Although note organization depended somewhat on the individual's note taking style, group notes were generally more organized than personal notes. Personal notes

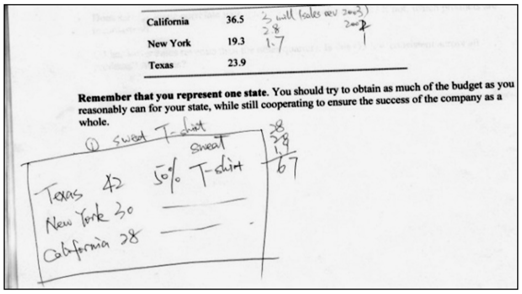

(a)

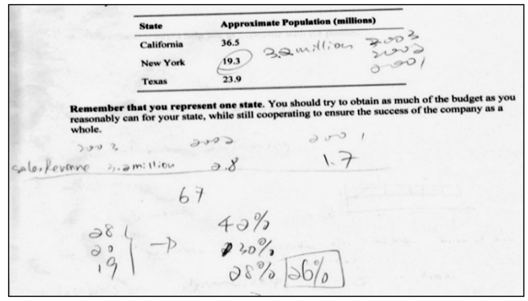

(b)

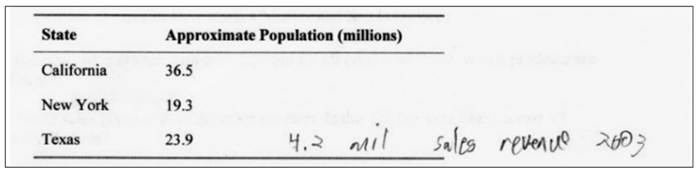

(c)

Figure 3.6: Three personal notes taken by participants in one group. This group did not have a designated note-taker.

were not always organized or written legibly or in a way that everybody could understand at a glance. Sometimes authors of personal notes used abbreviations or symbols that could be interpreted only by the note-taker. Possibly they were writing as fast as possible to minimize distraction, as taking notes was not their primary focus.

Notes' Usage Notes were used for a variety of different purposes, most commonly to further analyze findings and facilitate the problem-solving process (analysis phase), validate or remind the person of something (all phases), and create the final report (dissemination phase). Users referred to notes most often during the analysis phase. Saved values, calculated percentages, drawn charts, and other information helped users to make comparisons and reach decisions. Notes also facilitated the problem-solving process by recording the direction and sequence of the steps taken. This could help users to more easily determine the next step. For example, by recording the names of the charts created or values calculated or observed, participants could determine the completeness level of the task (e.g. what and how many more charts were to be created). Figure 3.7 shows an example of recorded analytic steps and the findings of each one. It was filled in gradually as information was found in various charts. The figure shows the completed version, indicating that participants had finished their calculations for all three states.

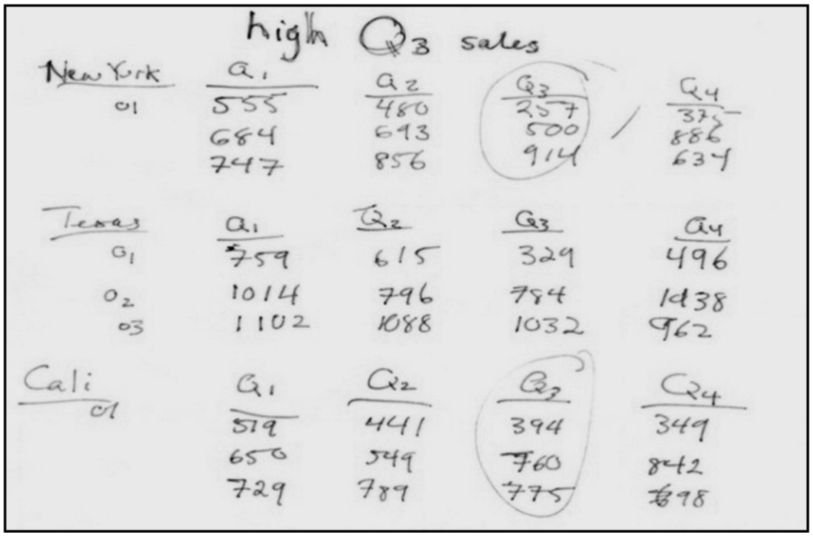

Figure 3.7: A participant organized information in a tabular format. This note shows that the group had calculated values for all four quarters and for all three states.

3.3.5 Awareness with Respect to Note Taking

I noticed that the manual note taking process impacted awareness among group members. This can be seen in Figures 3.1 and 3.8, where note-takers are disconnected from the other team members. In both cases, the note-taker has lost her direct view of the screen, either because she is forced to look away to see her notes or because other participants blocked her view. This role division is not necessarily unproductive, but it is possible that the tool design forced this work style, which might not be always desirable. Sometimes participants lost a sense of what others were doing while they were taking notes, and then had to catch up. For example, the person who took the note in Figure 3.5 was assigned the role of note-taker. He was sitting most of the time observing others (who were exploring data and creating visualizations), and therefore was unable to work directly with the application for much of the time. Although this division of roles may not have been unproductive for the group, it did deprive one group member of the opportunity to participate equally in analysis activities.

3.3.6 Wall Display Versus Tabletop Display

Table 3.1 shows a difference between record-keeping strategies/behavior among different types of displays. The first four groups used the tabletop, Groups 5–8 used the wall, and the ninth group used both. Only groups that used the wall display had

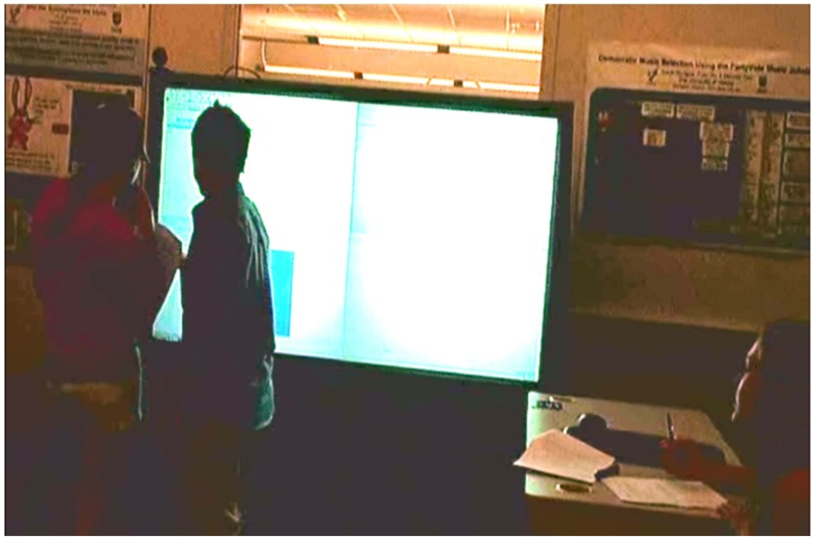

Figure 3.8: A seated note-taker of a group working on the wall display had her view obscured by two other participants.

a designated note-taker. This might be attributed to the different affordance of the displays; for example, fewer people can comfortably gather close enough to the wall display to interact with it than when using the tabletop.

In line with the literature [87], my interviews revealed that a wall display could support larger groups of people and provide a common view for presentations. On the other hand, the interactive tabletop display offers the potential for supporting formal and informal collaborative activities, such as planning, designing, and organizing. Most of the eight groups working on only one type of display mentioned that they preferred a wall display for audience-based situations such as presentations and a tabletop display for more collaborative situations. The ninth group, who used both displays, expressed a similar preference toward wall and tabletop displays

With regards to note taking, I observed that a note-taker of a group working on the wall display was more disconnected from the group (e.g. compare Figures 1 and 8). This was partly because the note-taker usually had to sit down or lean against a surface to take notes. In addition, often other users obscured the vertical screen by standing in front of it, so it was difficult for the note-taker to keep track of what was happening on the screen.

This was a less significant issue with the tabletop because the note-taker could stand side by side with others and take notes while holding the paper, as shown in Figure 3.1. Occasionally the note-taker did make use of a table positioned near the tabletop display. In these cases, the note-taker was not entirely disengaged from the rest of the group.

3.4 Discussion

In this Chapter I investigated RQ1 (How can collaborative record-keeping activities be characterized in the process of visual analytics for a small team of collocated collaborators?). This research question was open-ended but the importance of record-keeping emerged as a pivotal activity during this study. The observations showed that both taking notes and saving charts were frequently employed by participants to help them perform analytic tasks. Therefore, in the next section, I use evidence from this observational study to propose more specific design guidelines and considerations for visual analytics tools, with a focus on note taking. Results of this study are subject to some caveats. Later in Chapter 6 I discuss threats to validity and address the generalizability, validity and reliability of my results.

3.4.1 A Clear Need for Record-keeping Support

The main result is the importance of recording findings and cues (as notes or saved charts). Although some previous research [39, 43, 49] has suggested allowing annotation of visualizations, this study highlighted the importance of note taking as a critical activity. Taking and using notes was a frequent activity in all phases of the collaborative decision-making process. Lack of support for record-keeping had negative consequences such as disruption to workflow and decreased awareness of group activity. The importance and difficulty of record keeping was somewhat unexpected, as I did not tell participants that they should take notes or save charts and expected the major bottleneck to be interaction challenges with the single-user software. This highlights the need to build explicit record-keeping support into collaborative visualization tools. Recently, some research [30, 52, 97] has demonstrated how this can be done for single users, but work remains to extend this idea to multi-user systems.

3.4.2 Impact of Task Nature on Note Taking

This study further corroborated Heer and Agrawala's [36] proposition that the nature of the task affects both collaboration style and division of workspace. My observations revealed that Task 1, which involved focused questions, encouraged a highly coupled collaborative style of work, while Task 2, which required competition, led to a loosely coupled collaborative work style. In the interviews, most of the participants said that they would have preferred to explore information for Task 2 individually and then later share their results. As a result, notes taken in Task 1 were public whereas notes taken in Task 2 had a combination of public and private scopes. This finding emphasizes the need to support both individual and jointly coupled activities, as previously suggested [48, 50, 86, 94]. More importantly for us, it suggests that both group and individual record-keeping is necessary. An effective collaborative analytics system should provide both public and private records that are easy to distinguish and enable users to seamlessly switch between them.

3.5 Suggestions to Support Note Taking

How to best design effective record-keeping functionality for collocated work is not entirely clear. Here I offer some suggestions, which vary depending on the nature of the collaboration, whether or not the record involves data (or is linked to data), and whether the note is taken for group or personal use. Though these suggestions are based on the results obtained in this study, owing to use of students as participants, the number of participants, constraints of the software, and other limiting factors, research is needed to further assess these suggestions.

3.5.1 Integration Level for Notes and Saved Artifacts

Should notes be integrated with a history mechanism (i.e. along with saved artifacts and system states) or kept as a separate "notebook?" The analysis suggests that either answer would be too simplistic. Some notes, especially annotations and other notes of findings, have a clear link to an artifact that helped to form the insight. For example, a user might save a chart of revenue across different states, note that revenue is highest in New York, or write a reminder to later break down the New York revenue data by year and quarter. In these cases, the record is either a particular representation of data itself or can be linked to one. Retaining this link and enabling the user to return to the artifact and system state would have a clear benefit. However, notes also served as cues (e.g. reminders) or collected together findings from a variety of sources; in these situations, a notebook style is more relevant.

As a result, I suggest a hybrid model in which notes can be collected together in notebook pages but parts of a note could link to related artifacts, which might be stored in a chronological history. This would ensure that users could easily refer to source data when reviewing any given note. Artifacts could similarly link to the notes associated with them, and might also be directly annotated with drawings or text. I expect that such functionality would also simplify the task of recording findings. For instance, instead of writing down "Menswear has higher sales than womenswear," the user could simply circle the male bar in a bar chart showing sales broken down by gender.

Such notes could be captured on a shared display within specialized notebook containers that ideally would support both text and diagrams. Furthermore, these could be treated by the system as if they were artifacts such as charts. For example, they could potentially be added to a chronological history in the same way that a chart would be added. This would capture the development of the note over time, making it easier to understand the process that was followed. This might be particularly useful for helping a novice to learn the process that an expert analyst followed, or to help an analyst who is new to the project understand what work was done by previous analysts. Finally, because the number of notes and data artifacts can grow quite quickly, I believe that searching and filtering both types of objects will be important.

In the next chapters I look into this further by creating a system that links charts and notes. In Chapter 4 I discuss implementation of a very similar hybrid approach into CoSpaces (Collaborative Spaces). In CoSpaces, notes and visual snapshots are linked to the underlying analysis state so that the state can be easily reloaded. Notes are placed in the notes pane in chronological order. Keeping all of the notes in the same place makes the the reviewing process easier. The notes pane scrolls when the available space is exceeded. Users' feedback about this approach was very positive and they reported that it was very helpful for reviewing, validating and discussing previously found materials.

In Chapter 5, I address implementation of CLIP (Collaborative Intelligence Pad) for the context of intelligence analysis. CLIP also has a note pane that makes the reviewing process of notes easier. CLIP offers many options for searching, sorting and filtering notes based on the keyword assigned to the note (e.g. hypothesis, question, finding, reminder). CLIP takes this approach a step further by adding collaborative options, where an analyst can choose to integrate collaborators' notes into his/ her own view. Then the analyst can sort, and filter them based on keywords, time created, etc. To distinguish collaborators' notes easily, each collaborator has been assigned a distinguishable colour. Those colours reflect on any item created by a collaborator such as a note.

3.5.2 Notes for Group Versus Individual Use

Group notes could take the form of a shared history/notebook, plus shared note containers or papers as described above. To keep track of who did what, they might be spatially organized or colour coded by user. Individual notes present a greater problem because users may wish to keep their notes private or may want to avoid the burden of viewing all other users' notes. At the same time, individual notes occasionally need to be shared. One possibility is to provide private space within a shared display if there is sufficient screen real estate and if the notes are not confidential. Another alternative is to provide each user with a private display such as a tablet or digital paper. These could be linked to the common display to enable sharing. Ordinary paper notes or an unlinked private display are also viable options, but are more difficult to share with several people at once.

I was interested to note that the participants were less concerned about privacy of their individual notes and more concerned with the basic ability to work in parallel on individual tasks in preparation for discussion. This suggests that while individuals need their own space to work and to keep personal records, they may not need that space to be inaccessible to others. This confirms that shared display environments are reasonable for such tasks as long as individual workspace can be provided. Later in chapter 5 I observed that collocated collaborators preferred a fully integrated view. However, it should be noted that unlike this one the task in chapter 5 was non-competitive.

3.5.3 Record-keeping to Support Different Collaboration Styles

I observed tightly coupled work, where a shared history/ notebook would probably suffice. For loosely coupled work, participants may need to corroborate and combine the outcomes of their individual work. In this case, it may be better to give each individual personal space to work independently, but also to allow sharing. Individual notes and history items that could later be merged together could allow each user to track their individual work and then later compare it with the work of others.

Note that although they allow private work, individual desktops may not be the best solution here because they make sharing cumbersome.

Both CoSpaces and CLIP aim to support individual and group work and different collaboration styles. In this spirit, both systems provide an individual workspace that can be used for individual exploration and shared with group members later. CLIP also supports merging individual work for better collaboration and coordination of work (see Chapter 4 and 5 for more details).

3.5.4 Note Manipulation and Management

Note management and manipulation is also important to consider. Highlighting, grouping, and summarizing notes will increase their usability [10, 57, 97]. Previous research has shown the importance of providing users with functionality to manage and structure their notes. For example, in the context of education, it has been mentioned that a matrix structure for recording notes is more beneficial than a linear structure [7]. I also observed that most of the group notes were naturally taken in tabular format (Figure 3.5), which facilitated comparison.

3.5.5 Input Mechanisms to Support Note Taking on Shared Displays

Another important consideration is the form of input to be used for note taking in a shared display situation. I do not believe that there is one perfect solution, but in this section I hypothesize the trade-offs of several possibilities.

- **Wireless keyboard:** A wireless keyboard is familiar, easy to use, enables long notes to be quickly entered, and is not disruptive to other users. However, it may be difficult to place on or around the working area. In addition, more than one may be needed simultaneously and a good approach would be needed to associate each keyboard with the note a user wants it to effect. It could be effective for both tightly and loosely coupled collaborative work.

- **On-screen Touch keyboard:** This approach has the advantages that no external hardware is needed and the interaction concept is familiar. On the other hand, virtual keyboards are usually not very easy to use (and even more so on vertical surfaces). Like a physical keyboard, a virtual keyboard might be effective for both tightly and loosely coupled work styles. In chapter 4 for the second study, I used on-screen keyboard. However, participants found it awkward to use.

- **Stylus Input Plus Handwriting Recognition:** This approach may be easier to use than a virtual keyboard, provided that the user can rest his or her hand while writing (which is not possible with some touch surfaces). Possible disadvantages are low writing precision if the stylus is thicker than an ordinary pen, and difficulty of character recognition. As writing on a vertical surface (and sometimes also on a tabletop) could be awkward, external tablets might be employed for input and then linked to the shared display. Stylus input seems a good choice for annotating visualizations as well as taking notes in loosely and tightly coupled work. Isenberg et al. [49] investigated writing relatively low-resolution notes and annotations on high-resolution displays, but this is still a considerable challenge and research is ongoing.

- **Digital Paper:** Notes could be captured on digital paper and then directly entered into the shared display when desired. This is similar to the way that physical and digital objects are merged in systems such as the ColorTable [66]. The many benefits of tangible interaction are well known [79, 95, 103]. Using paper would enable a note-taker to edit notes without physically interfering with other work, allow the notes to be easily passed around, enable notes to be kept private, and enable flexible spatial layout and navigation of note pages.

- **Audio Recording:** Capturing notes by audio input is likely faster than other methods, and therefore may be less disruptive during tightly coupled work. Unfortunately, it is more difficult to use afterwards than written text, more difficult to search, and referring later to a recorded “voice note” is less likely to happen. Audio notes could be transcribed, but this is far from perfect with current algorithms. There is also a need for extra hardware. Audio notes are likely to be very disruptive to other users when working separately and would therefore only be useful for tight collaboration.

3.6 Conclusion

In this chapter to address RQ1, I characterized phases and activities involved in collaborative visual analytics for collocated groups. The results identified record-keeping as a process that is intensively used by data analysts but not well-supported by tools available at the time. In addition, I characterized notes according to whether they were findings or cues, and whether their scope was for personal or group use. I also described how notes were taken and used within four identified phases of data analysis. These analyses enabled me to offer many suggestions of how to better support record-keeping activities within visual analytics tools.

The result of this study led me to define my next research question:
How can record-keeping functionality be integrated into a collaborative VA tool in a way that supports visual analytics and collaborative mechanics? In the next Chapter I discuss the results of the second user study evaluating a collaborative VA tool that I designed to address this research question.

Chapter 4

CoSpaces-A System to Support Record-Keeping in Collocated Collaborative Visual Analytics

In Chapter 3, I discussed the results of my exploratory study which demonstrated that record-keeping is a critical component of the collocated collaborative VA process. However, this study was limited in that it used a system with no built-in record-keeping capabilities. Thus, although the study clearly showed that record-keeping was important, it could not assess whether best practices for single-user record-keeping support would extend to collocated collaboration. In addition, based on findings of this study I proposed a list of design suggestions to support record-keeping activities during collaborative VA. To better evaluate these speculations and better understand the design requirements for collaborative setting, I defined the second research question as: **How can record-keeping functionality be integrated into a collaborative VA tool in a way that supports collaborative mechanics?**

To this end, I designed CoSpaces (Collaborative Spaces). CoSpaces is a prototype tool for collocated collaborative VA on interactive tabletops that incorporates a visual record-keeping module.[1] I designed this tool based on the results of my previous study, described in Chapter 3, as well as available guidelines to design collocated collaborative VA tools. In addition, I organized a workshop with business intelligence design experts to brainstorm design ideas for this tool.

[1]This paper was a collaborative work published at the HICSS 2013 conference [64]. I would like to acknowledge my colleagues: Tyler Weeres for his help during the system implementation and Ali Sarvghad for his contribution to the design, observations and analysis of the results of this study.

There are many challenges and requirements unique to designing tools for collocated collaborative visual analytics [43, 93]. Different bodies of work have addressed some of these challenges; however, relatively few tools for collocated collaborative visual data analysis have been developed to support record-keeping activities. In Chapter 2, I discussed challenges most related to my main contributions such as supporting transient collaboration style, awareness, territoriality and supporting communication and coordination.

In this chapter, I describe the design of CoSpaces plus a user study that I conducted. This user study revealed that the design ideas employed in CoSpaces can effectively support group analysis tasks. In Chapter 2 section 2.3 I discussed Gotz et al.'s [29] categorization of various visual analytic behaviours. Their four-tier hierarchy is comprised of tasks, subtasks, actions and events. In this chapter, I report results of the user study that include a description of users' actions on recorded history items and notes and their foremost intention/s for those actions. While these actions can be conceptually placed within the action tier of Gotz et al.'s categorization [29] of VA activities, my categorization extends their work to collaborative work. Moreover, I report my observations of how these actions were distributed across different analysis phases (information foraging versus discussion) and collaboration styles (loose versus tight).

As with the previous study, my focus is on situations in which small groups of people gather face-to-face to perform visual analysis tasks and discuss their findings. Based on the observations, I suggest design considerations that would better tailor record-keeping modules for synchronous collaborative visual data analysis. The main contributions of this chapter are described in Chapter 1 as C4, C5 and C6.

4.1 Introduction

Record-keeping refers to activities that involve recording any information related to the analysis task. This recorded information can include recording visualization snapshots, system states, notes, and annotations. In Visual Analytics (VA), a record-keeping repository typically consists of recorded visualization snapshots, system states, notes, and annotations. For single users, record-keeping helps to generate insights and synthesize knowledge [37, 39]. While data analysis histories might be more important for collaborative tasks [108], the role and value of record-keeping in a collaborative context is less well studied and prior research has been mainly focused on asynchronous distributed VA. Moreover, general design guidelines for collocated collaborative VA tools [5,38,45,46], do not include much guidance for record-keeping. In this setting, record-keeping can facilitate building of common ground, sharing of visual artifacts and findings and facilitate offline discussions [39, 108].

Intrinsic attributes of collocated synchronous collaboration may impose different design requirements for record-keeping than in other settings. For instance, tabletops are well-suited to collaborative work in visual analytics since they allow users to interact and explore a dataset simultaneously. However, there are challenges to be considered while designing record-keeping modules for tabletops. In particular, I focus on the following three challenges: supporting transient collaboration styles (i.e. loosely coupled when collaborators work in parallel and closely coupled when working together), providing awareness, and enabling record-keeping and re-use.

Towards gaining a better understanding of how to design record-keeping, I designed and evaluated CoSpaces, a system for collocated collaborative Visual Analytics on interactive tabletops. With CoSpaces, I address the issues of territories for supporting parallel work on tabletops. CoSpaces extends the notion of a personal work territory with mechanisms to support awareness. In particular, it employs a history mechanism combined with a tab metaphor to enable users to review the work of others without disruption. In the design of CoSpaces, in addition to supporting record-keeping activities, I paid careful attention to salient issues of transient collaboration style, awareness, and territoriality. CoSpaces allows collaborators to smoothly transition between different collaboration styles, work individually when they needed while maintaining awareness of each other's work unobtrusively and easily share their results with each other.

In the collocated synchronous context, the closest research to my work is Cambiera [45], a tool that tracks each individual's history while they analyze a document corpus. Using colour-coding, past searches and documents are visually represented to increase users' awareness of each other's work. In the next chapter, I discuss how Cambiera is limited in term of providing support mainly for the early phase of the sensemaking process and how we should add to this work by investigating support needed for later phase of sensemaking (See Chapter 2 section 2.5.3).

Similar to Cambiera, in this chapter I investigate how access to a history of group members' actions influences awareness and discussion; however, I focus on tabular data, which has very different design constraints than a document corpus. In particular, I track users' created charts and notes rather than a record of which documents have been viewed. Somewhat less related is MemTable [41], a smart tabletop surface that captures and visually represents the table contents during meetings including individual participation histories. However, MemTable was designed for more general types of meetings rather than visual data analysis.

4.2 Overview of CoSpaces

To investigate visual record-keeping in a synchronous, collocated, collaborative setting, I designed and implemented CoSpaces, a visual data analysis tool with embedded visual record-keeping. Because there was no existing tool that incorporated current best practices for collocated collaborative work, record-keeping and tabular data visualization, I decided to build a tool rather than using an existing one that was not designed for this context. I designed CoSpaces for a large multi-touch tabletop display since such devices are thought to facilitate collocated work. The following subsections describe the primary features of this tool.

4.2.1 Worksheet

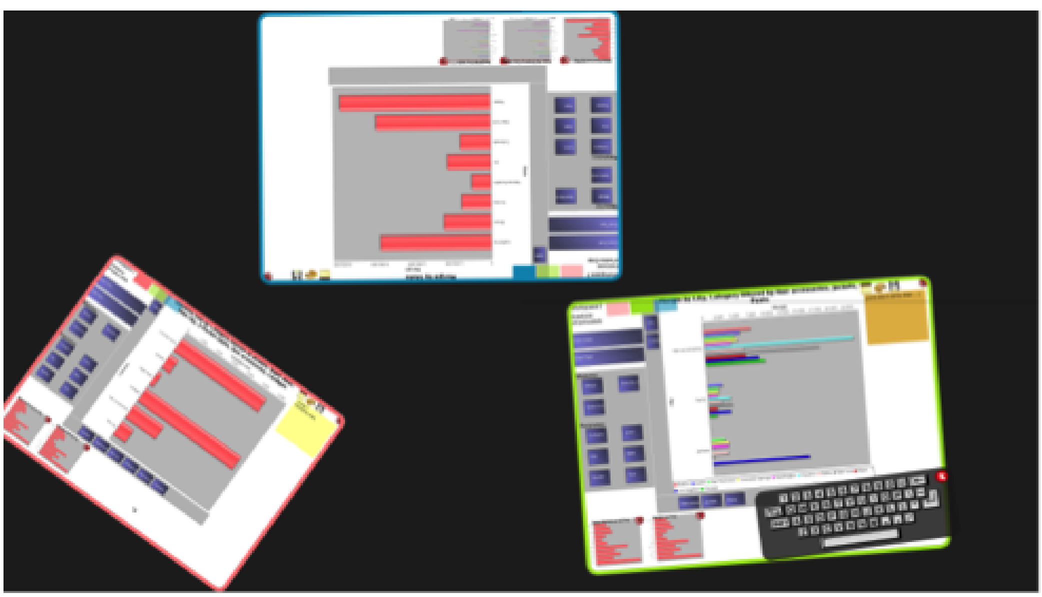

Figure 4.1: Dark background is the tabletop surface. There are three open Worksheets.

The CoSpaces interface is composed of Worksheets, as shown in Figure 4.1. The Worksheet was designed using the principle of “one space, many uses. Its design

Figure 4.2: Worksheet Details: Analysis pane (A) for creating and modifying charts, Visualization pane (B), History pane (C), Notes pane (D), and Tabs (E) that provide a portal view to other worksheets.

provides a team with the flexibility to work collectively on one or more Worksheets, or separately and simultaneously on multiple Worksheets. Each Worksheet defines a work territory, either personal or shared. Worksheets therefore enable both individual work territories and shared work territories, as advocated by Scott et al. [93]. Moreover, users may create several Worksheets, perhaps to investigate different data attributes and compare them side-by-side.

Personal versus shared Worksheets are identical as far as the system is concerned; ownership is defined by the way in which they are used. This makes it easy for users to convert a personal space into a shared space or vice versa. Worksheets can also be moved and resized. Each Worksheet's relatively wide border is uniquely coloured with a bright distinctive colour. This enables users to easily distinguish Worksheets from each other. Sections of a Worksheet are shown in Figure 4.2.

4.2.2 Tab Portal Views

CoSpaces uses a tab metaphor to facilitate awareness of other users' activities plus sharing of artifacts. Coloured tabs at the top of each Worksheet Figure 4.2E are associated with other existing Worksheets. Each tab is colour-coded to match the border colour of the Worksheet that it links to. Tabs act as portals to view other Worksheets. Tapping on a tab replaces the local worksheet content with a view of another Worksheet. Tapping on the local Worksheet tab switches the view back. When another Worksheet's tab is selected, the contents of all panes are changed to reflect the remote information, including the current visualization as well as recorded items in the history and notes panes. The user may browse charts and notes to learn about another user's past analytical activities and interests. To prevent unintentional changes and interruption, a Worksheet's remote view is read-only and navigation in a remote view is not linked to the other Worksheet's local view. To share charts, one can select an item in the history pane of a remote view and copy it to the local Worksheet's history pane.

I speculated that use of tabs for remote viewing could have two main use cases: (1) to gain awareness of others' work progress and direction during loosely coupled work, and (2) to review and share findings during closely coupled work. These speculations were supported by my observations, as I will show later.

4.2.3 Visual Record-Keeping

A critical part of the tab mechanism is the ability to see the past work done by others. The record-keeping module also serves to track and facilitate individual work. Analysts can review their previously created visualizations and reuse captured artifacts to perform analytical tasks such as chart comparison.

I designed record-keeping such that it facilitates capture of both visual artifacts (i.e. charts) and users' externalizations (i.e. notes). Notes and visual snapshots are linked to the underlying analysis state so that the state can be easily reloaded by tapping on a note or dragging a thumbnail to the central area. I define an analysis-state as the information that is required to replicate a system state at a later time (i.e. mapping and filtering information plus the chart type).

A Worksheet automatically captures and saves a copy of the current analysis-state right before a change, made by the user, has been applied. I use a simple heuristic inspired by the chunking rules of Heer et al. [37] to reduce history repository size.

An analysis-state is saved only when a change in the current mapping of data takes place. In other words, adding or removing filters will not result in a save. An analyst can externalize findings, hypotheses and so on using the notes pane. The importance of connecting externalized material to the visual representation of data has been previously recognized [14, 61]. Therefore, I automatically create a link between the current chart and the note.

As part of the analysis-state, I capture a thumbnail picture of the chart. Thumbnails are placed in the history pane in chronological order from oldest to newest (Figure 4.2C). The pane scrolls as the number of thumbnails grows. Notes are placed in the notes pane in chronological order, matching the chart thumbnails. The notes pane scrolls when the available space is exceeded. Moreover, since a note and its corresponding visualization are linked, an analyst can easily reload that specific visualization from the note.

4.2.4 Implementation

CoSpaces is multi-touch application written in JAVA. The Multi-touch for Java (MT4J) open source API was used to provide multi-touch capabilities within the Java code. CoSpaces uses the TUIO protocol to communicate with the finger and object tracking software, Community Core Vision (CCV). I used the JFreeChart chart library for creating and displaying statistical charts of data. I have adopted prevalent multi-touch gestures for scaling, rotation and translation. User interface objects such as Worksheets and on-screen keyboards are scaled, rotated, and translated by using two or more fingers. I have also implemented target highlighting to facilitate chart creation. As users drag a data dimension into the Visualization area, sections of the Visualization pane highlight only if the selected data dimension may be dropped in that location.

4.3 User Study

I observed pairs of participants working collaboratively on an analysis task using CoSpaces on an interactive tabletop. My goal was to gain a better understanding of how people use visual record-keeping in a collocated collaborative context and how visual record-keeping influences the collaborative analysis process. Therefore, I focus primarily on users' actions that involved history items and notes. Details about this study including consent form, introduction to the task and system, tasks, questionnaire, and follow up interview questions can be found in Appendix B.

4.3.1 Participants

I recruited 10 pairs of computer science students (16 graduate students, 4 undergraduates; 15 male, 5 female) who were familiar with basic data analysis activities and basic statistical charts. Age ranged from 19 to 35 (average = 27). Pairs were not required to know each other beforehand. Participants were compensated with $20 each.

4.3.2 Tasks, Dataset and Procedure

Prior to the user studies I run three pilot studies. This was mainly to ensure that 1) the analysis tasks were clear and understandable, 2) discover possible bugs in the prototype tool, 3) the introduction to task and prototype, follow up interview questions and other instructions were adequate and unambiguous. During these pilots, I also determined ideal times for the different steps.

Participants performed two tasks in which they could use system features freely and were not explicitly required to take notes or save charts. After a 20-minute introduction, they started a training task (Task 1), which took about 30 minutes and focused on learning CoSpaces (details can be found in Appendix B). They could ask either of the two observers if they had any questions.

After Task 1, each group was given a short 5-minute break to rest and read Task 2. Task 2, which took almost 40 minutes, was an open ended question. The following is a concise paraphrasing of Task 2: "Assume you are a financial analyst of a clothing company. There have been some unpredicted trends in the sales of most popular product lines that are: Dresses, Sweaters, Outerwear. You will look at the first three items and your team member will look at the rest. Analyze the sales data and at the end prepare a report for your CEO that explains possible reasons for the sales anomalies." The two tasks were followed by a questionnaire and a follow up interview that took almost 20 minutes.

Transitions between tightly and loosely coupled work happen naturally in collaboration [102]. In practice, participants might know the main analysis task in advance and perform some pre-analysis before the meeting. Later, during the meeting, sharing of their findings and insights would provoke further discussions and raise new questions. At this point, if the problem requires, the team may break up to investigate the newly posed problem, and join together again after finishing the ad hoc analysis. Task 2 tries to replicate as much of this process as possible, but leaves out the pre-meeting analysis for logistical reasons. The motivation behind Task 2 was to ensure I would observe use of record-keeping under both loosely and tightly coupled collaboration. The dataset used for this study was the same one used in the first study (described in Chapter 3). It include sales revenue, margin and quantity sold of clothing items in eight US states for three consecutive years, and consisted of 9 columns and 3273 rows.

4.3.3 Apparatus

For this study, I used a rear-projected 70-inch (diagonal) tabletop with a resolution of 3840 x 2160. The tabletop used a rear mount infrared camera to detect a (practically) unlimited number of touches.

4.3.4 Data Capture and Analysis

My colleague[2] and I independently observed users' interactions. I also videotaped each session. 400 minutes of video data were collected (around 40 minutes for each session). Then we manually coded the video data using a two-pass analysis approach. We first analyzed videos together to identify a set of repeated actions on history items and notes. In the second pass, we coded each individual's activities using the defined set of actions. My coding and qualitative observations are based on Task 2, as Task 1 was only intended as practice.

4.4 Findings

I first list the set of observed actions on recorded material, and the actions' relationships to analysis phases and collaboration styles. Then I discuss emergent record-keeping behaviours in the collocated collaborative context. I will also report observed instances of visual record-keeping use with no direct interaction with record-keeping mechanism; that is, cases where users clearly looked at a recorded item but did not touch it. I name these soft actions.

[2]Please note that this research was result of collaboration with my colleague, Ali Sarvghad. Thus any referral to "we" in this section acknowledges his work.

All the reported observations are based on Task 2, as Task 1 was only intended as practice. I focus on qualitative observations and participants' comments from the interviews, and structure the discussion around the major components of CoSpaces. However, for completeness, I also include quantitative results from the questionnaires and qualitative results from the observations and videos.

4.4.1 Actions on History

I identified nine primary actions (Table 1). These can be conceptually placed within action tier in Gotz et al.'s categorization [29] of VA activities. Similar to Gotz et al action-layer, the actions derived from this study are domain independent and carry some meaningful information regarding users' intentions (Figure 4.3).

Action on History	Primary Intention	
Reload Chart	Review Existing Chart	112
	Reuse Existing Chart	43
Review History	Review Charts	102
	Looking for a chart	26
Manual Save	Managing History	102
Delete	Managing History	99
Note Taking	Record Personal Finding/Insight	74
	Record Discussion Outcome	3
Review Note	Reload a Chart Linked to a Note	18
	Revisit all Notes	5
	Use Notes to Support Discussion	5
	Revisit a Note	2
View Current Chart	View Current Chart	11
Copy Local	Transfer a Chart between worksh..	4
Create External Worksheet	Comparison between charts	3

Figure 4.3: Actions on history and their primary user intentions. Right column presents the frequency of each combination.

Some actions such as review note and review history were done to achieve more than one analytic goal. To infer primary analytic intentions, I relied on my observations and action sequences. For actions that appeared to happen for more than one

Actions	Description	No.
Reload a Chart	Reload a previously created chart from the history, either the local history or a collaborator's history.	155
Review History	Review charts within the history, either the local history or a collaborator's history.	128
Manual save	Manually save a chart into the history pane.	102
Delete	Delete a chart from the history pane.	99
Note taking	Write down notes in the note pane.	77
Review Notes	Review note(s), add to a note, and / or reload a chart linked to a note.	30
View current chart	View the collaborator's current visualization through tabs.	11
Copy Local	Copy a chart from a collaborator's worksheet to the local Worksheet.	4
Create External Worksheet	Creating a new Worksheet using a chart from the history.	3

Table 4.1: Primary actions on visual record-keeping and the frequency of each.

reason, I analyzed pairs of action units. For instance, I observed that in most cases, a Review History action was followed by Reload Chart or another Review History action. Therefore, I concluded that primary user intentions for Review History were to look for a specific chart and to review a number of charts.

I recognize that my inferences may not always be correct, and so these numbers should be taken as approximate. For example, some instances of the Reload Chart action could have been to replace a wrongly reloaded chart. In addition, the frequencies of actions and primary user intentions are influenced by the system design and the individuals. For example, I suspect that there would have been fewer Delete Chart actions if the tool had used a smarter automatic saving algorithm that saved fewer charts. Nonetheless, I believe that the frequency of such misclassifications is very small, so that the high level trends are still valid or at least representative of the users using these tools.

4.4.2 Actions and Analysis Phases

I observed parsing, information foraging and discussion phases. At the beginning, participants in all the groups quickly discussed the problem and devised a strategy. All the groups then proceeded to an information foraging phase and started analyzing data and looking for and recording findings. This was followed by a discussion phase that involved sharing insights. I observed use of visual record-keeping under information foraging and discussion phases. Therefore this report of actions is limited to these two phases.

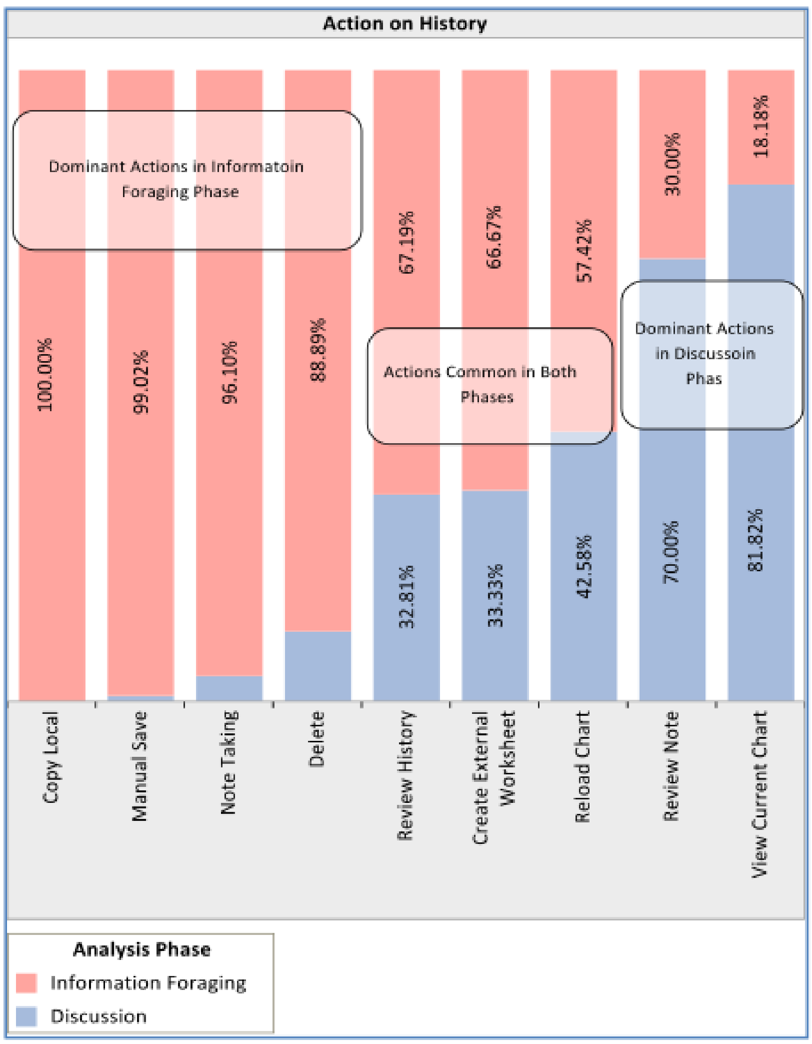

Figure 4.4: Record-keeping actions in each phase.

I distinguished phases based on the main analytical objective. During information foraging, participants gathered insight through visual exploration. They created charts and recorded their findings. During discussion, participants tried to form explanations and hypotheses around their findings. Because of the task design, all participants started with information foraging and then proceeded to discussion.

On average, groups spent 75% of the time on information foraging and 25% on discussion. Therefore the number of actions performed during information foraging (455) is almost triple that during discussion (154). Nonetheless, a closer examination of actions performed in each phase shows that there is a relationship between analysis phases and the actions performed on history. As shown in Figure 4.4, copy local, manual save, note taking and delete actions almost entirely happened during information foraging. Note reuse and view current chart were mostly performed in discussion. Reload a chart, review history and create external worksheet actions were common in both phases.

	Action on History / Primary Intention			
	Reload Chart		Review History	
Analysis Phase	Review Existing Chart	Reuse Existing Chart	Review Charts	Looking for a chart
Discussion	62	4	26	16
Information Foraging	50	39	76	10

Figure 4.5: User intentions for reload chart and review history actions in different analysis phases. Numbers show count of observed actions.

Of the actions that occurred frequently in both phases, Reload Chart and Review History were the most interesting because users' intentions for undertaking them changed depending on the phase (Figure 4.5). For example, during discussion, reloading a chart was much more likely to be for the purpose of reviewing it than reusing it for new analysis, whereas during information foraging, both purposes were common. Also, reviewing the set of previously created charts was much more common during information foraging than discussion, whereas looking for a specific chart was more common during discussion.

4.4.3 Actions and Collaboration Styles

The design of Task 2 required an independent, parallel work style, as defined by Isenberg et al. [46]. In this collaboration strategy, each participant works on the problem independently and results are then combined and discussed at the end.

As expected, I observed a major shift in collaboration style from loosely coupled to tightly coupled work as participants moved from information foraging to discussion and the goal changed from problem-solving to decision-making. Because of this high correlation, I cannot distinguish whether differences in actions on the history were caused by different collaboration styles, different high-level analysis goals, or both.

However, I note that brief changes in collaboration style did happen within phases. During information foraging, a participant sometimes stopped his/her collaborator to have a brief conversation. The content of these conversations was nearly always data analysis oriented; hence I classified these activities as tight collaboration but belonging to the information foraging phase. For instance, one participant stopped his collaborator to inform him of an interesting finding and asked him investigate the same data dimension for other patterns. Similarly, during the discussion phase, participants occasionally worked independently to review their work or do ad hoc analysis to validate a hypothesis. These short periods of individual work were almost entirely in line with the current decision making goal, and were classified as loose collaboration during the discussion phase. Although changes of collaboration style within a phase were infrequent (96% and 5% loosely coupled work in Information foraging and Discussion phases), they seem to be an integral part of collaborative data analysis dynamics.

I wondered whether using recorded material might invoke a change in collaboration style. For example, viewing one's collaborator's history and notes could inspire a conversation. However, the results indicated that this was not the case. I recorded the collaboration style during history use and immediately following the history use, and these were identical over 99% of the time (i.e. in all but one case).

4.4.4 Record-Keeping Behaviours

Participants used visual record-keeping to capture what they found important in the data, and their explanations for those findings. Based on my observations, I identified two dominant record-keeping strategies. The first was a note taking oriented strategy, in which participants took considerably more notes. The second was chart saving oriented, in which participants manually saved a larger number of charts. As shown in Figure 4.6, with the exception of groups 1 and 7, the rest of the groups exhibited primarily one of these behaviours. It is also evident from the same figure that capturing important findings predominantly took place during the information foraging phase. This makes sense since during this phase participants were investigating data for findings. In the discussion phase only a few notes were taken to record group discussion outcomes.

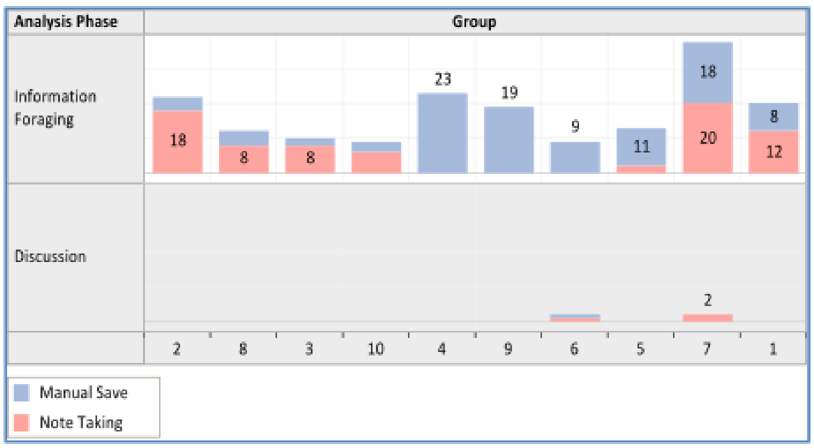

Figure 4.6: Count of Manual Save and Note Taking actions by each group. Groups are sorted by strategy.

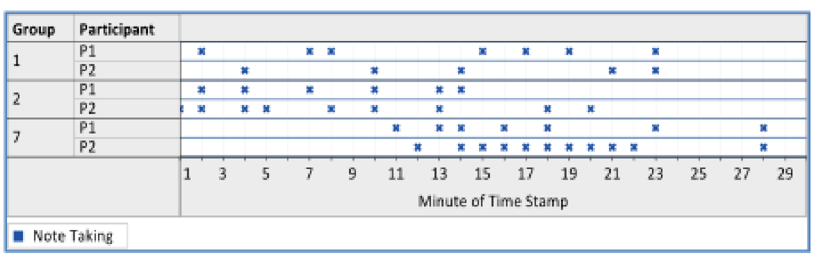

Figure 4.7: Note taking actions by participants in groups 1, 2 and 7.

Since participants simultaneously shared the surface, they had a high level of awareness of each others' interactions. Consequently, this could have influenced participants' analytic behaviour. For example, opening an on screen keyboard for taking a note by one participant could have been easily viewed by his/her collaborator. This in turn could have provoked a similar action by the other participant. A closer investigation of groups 1, 2, and 7 (Figure 4.7), who took noticeably more notes, shows that in many cases, a note taking action by one participant was closely followed by the other one. I hypothesized that one participant's behaviour influenced the other in these cases; however, this observation could have been coincidental and requires further investigation.

4.4.5 Use of Tabs

Tabs were used to view contents of another worksheet and were used almost equally in both information foraging (15) and discussion phases (17). During information foraging, tab use mostly took place in the middle of the phase, whereas tab use during discussion was almost evenly distributed throughout the phase. This observation matches the logical flow and objective of the phases. Participants started the information foraging phase by analyzing data and gathering findings and there was not enough history built to motivate remote viewing of a collaborator's work. As the work in this phase progressed and some history was built up, participants started examining their collaborators' work. During discussion, participants wanted to share findings and insights; therefore use of tabs to access past items happened throughout the phase.

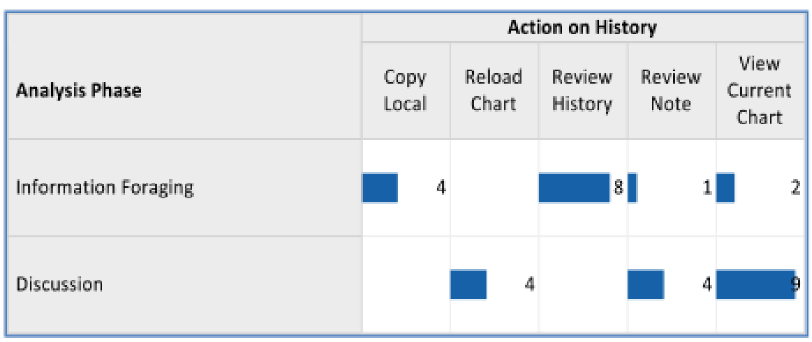

Figure 4.8: Actions on history while using tabs to view work of the other participant.

Actions on history during tab use were noticeably different between phases (Figure 4.8). This figure shows the count for actions taken by each participant while remotely reviewing collaborator's work. In information foraging, participants mostly used tabs to review their collaborator's work history (Review History, 8), reuse an interesting chart (Copy Local, 4) and observe the collaborator's current work (View Current Chart, 2). During the discussion phase, participants mostly used tabs to share a view (view current chart, 9), share findings (Review Note, 4) and review work (Reload Chart, 4). This dissimilarity can be attributed to the different objectives of each phase. Note that the overall counts are small, so relative proportions of events should not be taken as representative.

4.4.6 Quick Review

All the actions on history required physical interaction with the visual record-keeping module. In addition, I observed another, quite prevalent action on history that did not involve any direct physical interaction with the system (soft actions).

On several occasions, I observed head movement, suggesting that a participant quickly glanced at the visible portion of the history pane where the most recent charts were placed. This quick review happened under various circumstances. For instance, a quick review happened after almost every work interruption during the information foraging phase. It also often occurred before making a new chart. This could have helped participants to stay focused on the recent analysis path or to confirm that they had not made that chart already. Quick review could have also been performed for making non-detailed comparisons between the current visualization and charts in the history pane.

Without eyetracking data, counts of these quick review actions would be unreliable. I therefore only counted and categorized concrete actions on history (when there was a clear physical direct interaction with the system), and do not have quantitative information for quick review actions. Nonetheless, the observation that these quick review actions occurred suggests that visible thumbnails of recent visualizations provide useful support for data analysis.

4.4.7 Feedback on CoSpaces' Features

Figure 4.9 depicts participants' rankings of system features and their overall impression. Participants were generally positive about most design ideas introduced in Co-Spaces, but were slightly less enthusiastic about note taking because of the difficulties of taking notes using the on-screen keyboard.

Participants' quantitative assessments of the usefulness of the tabs were positive. Out of 20 participants, 17 assessed Tabs as useful in their evaluation. The average score given to Tabs was 4.95 with a STDEV of 1.07.

In the follow-up interview, 11 participants explicitly regarded record-keeping as one of the most useful features of the system. For instance one participant expressed

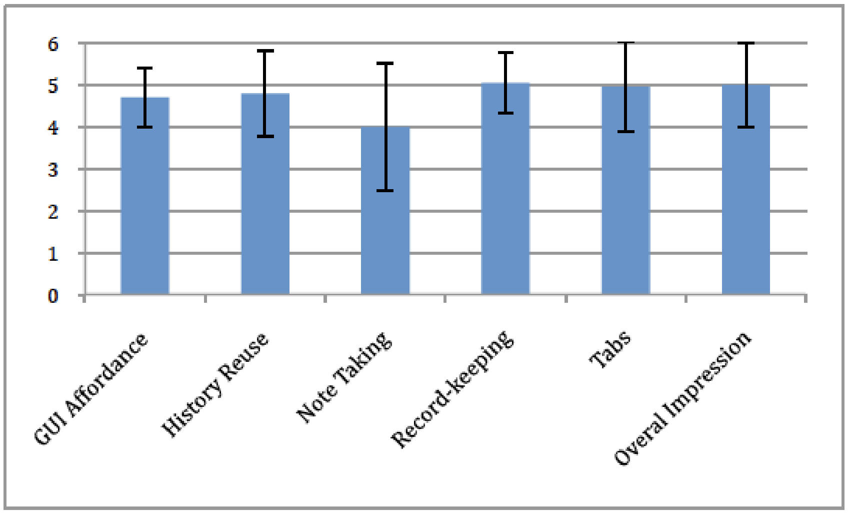

Figure 4.9: Quantitative questionnaire results showing participants' average assessments of CoSpaces. Participants provided ratings on a 6-point scale, where scores above 3 were positive and scores of 3 or lower were negative. Error bars show +/- 1 SD.

that "the ability to save the charts is great" and another one said "reloading from history was really fast and efficient and worked well for me."

Although my expectation of the current note taking mechanism in this version of CoSpaces was not very high, I still received a relatively large amount of positive feedback. The current note taking implementation was limited to the use of an on-screen keyboard for writing on simple post-it-like notes. The simple on-screen keyboard was straightforward to use, but was somewhat slow and awkward as an input mechanism, which is most likely why the note taking feature received a lower average rating than other aspects of the design. Nonetheless, I observed that many groups used the on-screen keyboard to take notes (total of 71 times, used by 8 out of 10 groups).

Participants also found the simple design of the Worksheet user interface effective and easy to use (Figure 4.9, GUI affordance). Based on these quantitative results, there was a consensus among all the users towards usefulness of these features. Note taking was comparatively less well-liked, but I suspect this could be improved by replacing the on-screen keyboard with a less cumbersome input method.

4.4.8 CoSpaces' Support for Different Collaboration Styles and Awareness

The flexible design of the Worksheet accommodated changes in collaboration style and adapted to versatile analytical behaviours. Figure 4.10 shows a collage of several groups in different collaboration styles. The first column shows group members working individually, and the next two columns show how participants discussed the data while referring to each other's Worksheet.

Figure 4.10: Collage demonstrating the CoSpaces support for changes in collaboration style.

Additionally, I observed that the design of the Worksheet not only supported changes in collaboration style but also could facilitate the analytical reasoning process. The ability to create multiple new Worksheets (Figure 4.11) as well as creating a

Worksheet from an item in the history pane facilitated exploratory analysis as well as comparison of charts. An example of this recurrent behaviour was group 6 (Figure 4.11, left segment). Both of the participants created new Worksheets to view their team member's work progress and put them side by side to compare two charts created for California and New York. When asked about this analytical behaviour in the follow up interview, one of the participants gave the following comment: "It was like working with actual paper and I liked the ability that you could put a workspace like a piece of paper aside when you are done with that". Group 8 (Figure 4.11, right segment) devised an interesting strategy for Task 2. At the beginning, they decided to create three Worksheets each. I asked the reason for this decision in the interview and they said that since there were three main data dimensions that they wanted to analyze, they decided to create a Worksheet to explore each dimension separately.

Figure 4.11: Working with more than one worksheet simultaneously.

My qualitative observations corroborated the usefulness of tabs for providing awareness. Many participants provided positive feedback about being able to view each other's work progress via tabs. For instance, in the follow up interview a participant expressed that "real time update of [the] other's view was interesting, because [I] could keep [myself] updated all the time". Another participant mentioned that "being able to see other's workspaces, [and] keep track of them in own workspace" was one of the most useful features of the system.

Because the groups contained only 2 people, and they often positioned themselves side by side (8 out of 10 groups), participants could easily look over at each other's Worksheets to see what the other person was working on. Participants frequently did this for an update on current work. However, I observed that while working individually, none of the participants attempted to get close enough to the other's Worksheet to have a detailed review of his/her work history; for this purpose, they used tabs instead. This behaviour avoids unnecessary interruption and imposition on the personal territory of another person. Thus, while tabs are useful for observing another user's current work, they may be even more beneficial for reviewing a collaborator's past activities.

4.4.9 CoSpaces Support for Note Taking & Reuse

Taking notes was one of the dominant record-keeping strategies. Users took notes to record their important findings during information foraging and also to document the outcomes of discussions. Most of the note taking took place during information foraging (74 out of 83) and note review mostly occurred during discussion (21 out of 24). With a total of 107 instances, note taking and reuse are two of the prevalent record-keeping actions, strongly suggesting that VA systems should support note taking.

Participants reported that the link between a note and its related chart was very useful. The fact that notes were linked to the charts was extremely important to participants. Seven participants explicitly mentioned the linking as important.

4.5 Discussion

To summarize, in this chapter I investigated RQ2 (How can record-keeping functionality be integrated into a collaborative VA tool in a way that supports collaborative mechanics?). I found that record-keeping played an important role in both information foraging and discussion phases of collaborative work, but that the types of actions and the reasons for them differed between the phases. During loosely coupled information foraging, users primarily used record-keeping tools to record their own findings, and maintain some awareness of collaborators' activities. During discussion participants primarily used record-keeping to share past charts and notes. Remote copies of worksheets were similarly useful during both phases but for different purposes (awareness during loosely coupled work and seeing the same view or finding charts during closely coupled work). I also identified two different strategies for capturing findings: chart-focused versus note-focused approaches. The chart-focused strategy required a lot of history management to keep a clean history pane.

The results demonstrate that in collocated collaboration, record-keeping tools do indeed play an additional role compared to their use in single-user systems. In particular, reviewing another users' recorded items through tabs was used to gain awareness and share work, and recorded materials served as a starting point for discussion. However, it was interesting to note that using recorded materials did not apparently lead to changes in collaboration style (e.g. from loose to tight).

The list of actions on history is based on my own observations and could be influenced by this study and tool design. For example, the limited history representation did not provide search and filter actions on recorded artifacts. Although these actions can be considered as special cases of review history, users' key intentions for performing these actions are probably different than a detailed review. Therefore, they could have been considered as independent entries in the list of actions.

The frequencies of actions and temporal distribution that I observed are also undoubtedly related to particulars of the study and design of CoSpaces. I suspect that the actions and intentions themselves would be repeated in other VA situations, but that their distribution over time and their relative frequency could change. For instance, with a group of three or more participants, I speculate that there may be more instances of tab use (or other mechanisms that provide unobtrusive remote access in a tool) to review a collaborator's history, since it would be more difficult to keep track of what everyone is doing. Similarly, a more complex task might lead to the use of more worksheets from history items in order to branch the analysis to a greater degree.

4.6 Design Implications

A prerequisite for carrying out actions on record-keeping is proper tool support. There have been general design guidelines for collocated collaborative visual data analysis tools [33, 36, 46], but no specific guidelines to support record-keeping in collocated collaborative VA. Here I suggest some design considerations to improve record-keeping functionality in collocated collaborative VA on large interactive surfaces.

4.6.1 Multiple History Views

Participants often applied a set of filters to different charts. A smart history representation algorithm could show history items with similar filtering criteria. This would facilitate history review when a user wants to find similar charts.

I also noted that participants remotely viewed their collaborator's history pane to acquire and maintain awareness. A record-keeping module could expedite this activity by automatically setting the default view to bookmarked (i.e. manually saved) items rather than including all automatically saved items, to reveal items that were found important by another user. Building common ground would be easier when collaborators share the knowledge of what is perceived as important by the others.

Participants in this study copied items from their collaborator's work to use in their own analysis. To facilitate this process, record-keeping could automatically identify collaborators' records that are related to a user's current line of inquiry. This recommendation could also increase awareness. In the next chapter (5) I discuss the automatic discovery and linking approach.

In addition to automatic representation of items in history, users should be able to optionally select different history views with varying levels of granularity. I suggest detailed, bookmarked, filtered and customized views. The detailed view would consist of all the saved artifacts by both the system and user, enabling a full review of the analysis path. The bookmarked view would only show items explicitly saved by the user. This view would be especially valuable during the discussion phase when group members are sharing their important findings. My observation that users preferred to manually save rather than manage automatically saved items suggests that the bookmarked view should be the default and that a more detailed view should be hidden until requested. A filtered view would present the results of searching and/or filtering the history items. Finally, a customized view could be comprised of items that are grouped/ordered based on a user's analytical needs. During the study, a few participants asked if they could arrange and group items in the visual history, suggesting that customization is important. In all these cases, a user should be able to easily select and switch between different views.

4.6.2 Support for Sharing

The results suggest that changes in collaboration style should be considered in record-keeping design. Participants used tabs under both loose (15) and tight (17) collaboration styles (Section 4.4.5). Based on this observation, I define *direct* and *indirect* sharing to distinguish between explicit and implicit sharing of work history. In the discussion phase, participants directly shared work to support decision-making. While working individually, participants used tabs to indirectly access their collaborators' work. Therefore, design of an indirect sharing channel should provide remote, unobtrusive and non-interruptive access to a collaborator's work.

One limitation of CoSpaces was that there were no privacy controls. Although this was not a problem in this laboratory study, I anticipate that it could be a concern in real world situations. For example, my first study (Chapter 3) showed that business analysts in a competitive scenario may wish to carefully guard what they share. Therefore, control over privacy levels and degree of sharing may be a good idea. Users should be able to change the visibility scope of an item among members of the group.

4.6.3 Support for History Management

Some participants frequently performed manual save and delete actions on the history. At least one participant reported that having to manually delete unwanted items was much more cumbersome than saving desired ones. This behaviour indicates the perceived importance of history as a container of important analytical artifacts. The observation of abundant "quick reviews throughout analysis also strengthens this speculation. These results suggest that manually saving items to the history may be a better choice in many circumstances over automatically saving them. Alternatively, an improved algorithm could be devised that would save the analysis state less often and better infer the states that are important to keep. A customizable view (as mentioned in 6.2.1) could also facilitate history management.

4.6.4 Support for Note Taking & Reuse

Taking notes was one of the dominant record-keeping strategies. Users took notes to record their important findings during information foraging and also to document the outcomes of discussions. To further facilitate note reuse, a record-keeping module could automatically create links between related notes based on their content. This would create a network of recorded insights and findings instead of individual notes, which would help analysts to "connect the dots, as suggested by [98]. In the next chapter (5) I address automatic linking related contents and its effects on collaboration in more details.

4.7 Conclusion

In this chapter, I addressed the second research question. I designed CoSpaces and evaluated the effects of incorporating a record-keeping module into a collaborative VA tool. I conducted a user study which corroborated the importance of supporting record-keeping activities in collaborative VA and extended this knowledge for collocated settings. CoSpaces introduced the concept of tab portal views that help to address the challenge of awareness, especially during periods of loosely coupled work. This user study indicated that CoSpaces includes some useful design concepts for collocated collaborative work on interactive tabletops. Most importantly, the tabbed worksheet design (with the incorporated record-keeping functionalities) enabled users to review the work of a partner without disruption and supported awareness during the loosely coupled work. In addition, the flexible uses and positioning of worksheets allowed users to seamlessly transition between closely and loosely coupled work styles.

Moreover, I categorized nine primary actions on history and key user intentions for each action. These actions and intentions varied depending on the analysis phase and collaboration style. While these actions can be conceptually placed within previously categorization of VA activities [29], my categorization extends previous work in the context of single user to collaborative work. In addition, I report my observations of how these actions were distributed across different analysis phases (information foraging versus discussion) and collaboration styles (loose versus tight). During information foraging, when group work was loosely coupled, history and notes were primarily used to record findings, support individual analysis, and maintain awareness of others' activities. During discussion, when group work was closely coupled, record-keeping tools were primarily used to help present past findings to collaborators and to record discussion results. Interestingly, a remote view of another workspace was useful in both situations, but for different purposes.

Based on these findings, I suggested providing various views of recorded material, showing manually saved rather than automatically saved items by default, enabling people to review collaborators' work in an unobtrusive way, and developing automatic algorithms to better identify which items are related and important to a user's current analysis task. In the next chapter, I further discuss design requirements to integrate record-keeping into VA tools. More specifically, I address RQ3:

What are the effects of linking collaborators' related externalizations to one another on awareness, performance, communication and coordination?

Chapter 5

CLIP- A System to Support Communication and Coordination in Collaborative Sensemaking

In Chapter 4 I reported on the design of CoSpaces, a prototype tool for collocated collaborative VA, and discussed the study and its results. One of the main objectives of CoSpaces was to support record-keeping activities. In order to support smooth collaboration, I paid special attention to other design guidelines for collaborative VA tools. For instance, I aimed to support transitioning in collaboration styles as well as awareness when collaborators want to switch from closely coupled work to loosely coupled work. Therefore, I designed tabs as an awareness channel. However, I observed that visiting each other's work through tabs was cumbersome. In addition, some of the findings suggested that having a more integrated view with automatic linking of investigators' findings might be a better alternative to consider. To move to the next phase, I conducted a literature survey to better understand other requirements to support collaborative sensemaking. In Chapter 2 I reviewed collaborative sensemaking processes and model and summarized the existing guidelines on how to support collaborative VA tools. Many tools have been developed to support both phases of sensemaking but with different focus. For instance, there were tools that support sensemaking process but that were designed for individual work (e.g. [6, 40, 116]). In the context of collaborative sensemaking much less has been done. Most of the existing tools either focus on the information foraging loop [45] or asynchronous collaboration [13, 115] with a different set of requirements as discussed in Chapter 2.

Therefore I was well motivated to design a new tool for collaborative sensemaking to support the sensemaking loop. While the design of CLIP draws upon prior research on sensemaking and collaboration support, its primary focus is on investigating other awareness channels to facilitate collaboration. I defined the next research question as **What are the effects of linking collaborators' related externalizations to one another on awareness, performance, communication and coordination?**

Automatic linking of collaborator's externalizations (aiming to reduce the barriers to seeing relevant work of collaborators) seems a reasonable approach to increase awareness in an unobtrusive way. In this chapter, I address the design of CLIP (Collaborative Intelligence Pad) and the user study and its results.[1] The main contributions of this study are introducing and evaluating Linked Common Work (LCW) technique and the effect of LCW on improving awareness, communication and coordination among team members (described in chapter 1 as C7, C8 and C9).

5.1 Introduction

Supporting collaborative sensemaking has been identified as one the unique challenges in collaborative visualization [44]. Sensemaking in collaborative visual analytics is a very time consuming and demanding process. It requires the team of analysts to iteratively exchange and discuss their results to be able to form and evaluate hypotheses, derive conclusions, and publish their findings. Team members also need to maintain *awareness* of each other's work, including both activities that people are working on as well as results and evidence that they have found. Tools that provide *externalization* support (i.e., ability to record insights, questions, and findings, for example as text notes) can help team members to record, organize and share their results [9, 38, 47, 115], and those that provide awareness channels should enhance collaboration, communication and coordination [26]. However, to date, we have a limited understanding of how to provide externalization and awareness support for collocated collaborative teams. How should such tool support look and behave within visual analytics tools?

I extended the use of LCW to facilitate synchronous collocated collaborative sensemaking and evaluated its effect on collaboration mechanics. With LCW, common work elements such as similar findings are automatically discovered, linked, and visu-

[1]The material presented in this chapter will appear in IEEE Transactions on Visualization and Computer Graphics [65].

ally shared among the group. I built this technique within a "collaborative thinking space" to enable analysts to record, organize, schematize and share their externalizations. Linked common work 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. The 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.

This work focuses on supporting teams of investigative analysts, for example in the domain of intelligence analysis. Intelligence analysts need to sift through large document collections, determine which pieces of data are relevant, and gradually build up an explanation supported by evidence. Field studies have revealed that professional analysts need to share sources and data, view each other's work, and combine findings together in order to build common ground, resolve conflicts, and validate each other's findings or hypotheses [15, 55].

The *sensemaking* process of intelligence analysts has been studied in some depth, and has been described as involving two iterative loops: the information foraging loop and the sensemaking loop [84]. The information foraging loop involves searching for relevant data, reading, filtering, and extracting information, whereas the sensemaking loop involves iteratively developing a mental model, forming and evaluating hypotheses, and publishing the results. I focus primarily on supporting later stages of the sensemaking process (i.e., the sensemaking loop), when teams are more likely to work together in a synchronous, collocated fashion. Kang et al. [55] found that intelligence analysts often choose to work asynchronously during information foraging, but typically work in a collocated synchronous fashion to synthesize their findings. This synthesis phase is reported to be the most difficult and time-consuming phase of analysis [55].

In this chapter, I explore the design of visual thinking spaces that support the sensemaking loop in collaborative analytics. A collaborative thinking space should enable analysts to record and organize findings, evidence, and hypotheses; moreover, it should facilitate the process of sharing findings amongst collaborators, to minimize redundant work and help investigators identify relationships and build a shared understanding. In this chapter, I discuss the value of employing LCW to relate and integrate team members' visual thinking spaces. The notion of LCW closely resembles *collaborative brushing and linking* [45] in which certain 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 activities. In contrast, my work facilitates later stages of the collaborative sensemaking process (i.e., the sensemaking loop), by applying the linking concept to people's externalizations (i.e., recorded findings and notes). I anticipate that enabling analysts to see how their findings relate to each other should make it easier to maintain awareness of each others' work, build common ground, and solve analytic problems. I address the following sub-questions related to RQ3:

- RQ3a: Does linking collaborators' externalizations lead to better analytic outcomes?
- RQ3b: Does linking collaborators' externalizations improve communication?
- RQ3c: Does linking collaborators' externalizations help collaborators to coordinate their work more effectively?
- RQ3d: Does linking collaborators' externalizations increase collaborators' awareness of each others' findings and activities?

To answer these questions, I designed and implemented CLIP. CLIP is a collaborative visual thinking space to support collaborative sensemaking. CLIP allows analysts to record their findings in the form of a node-link graph and timeline, add evidence to facilitate evidence marshaling, and add free form text to record hypotheses, questions, to-do-lists, etc. Most importantly, CLIP incorporates LCW to relate and integrate the findings of different collaborators. I assessed the value of LCW for visual thinking spaces by comparing CLIP to a baseline tool (BT) without the LCW features. Results of this user study demonstrated that LCW led to more effective group coordination and communication as well as better analytic outcomes.

As mentioned earlier in Chapter 1, the application domain that I investigate in this phase is intelligence analysis. I switched to intelligence analysis because I needed a dataset to evaluate the idea of LCW and VAST 2006 dataset had ground truth, so I could use that to test how well people actually did at sensemaking. Moreover, intelligence analysis has similar attributes to business intelligence in that both domains include a set of systematic and analytical processes in order to explore complex datasets, find trends, patterns and correlations to solve a problem. In addition, co-located collaboration is a frequent use case for this domain [18, 46, 55, 83, 84, 109]. However, it should be noted that the application areas are not the main focus of the thesis but rather the main focus is on understanding the collaborative VA process to inform the VA community in general.

5.2 Linked Common Work (LCW)

The LCW technique employed in CLIP reveals similarities between collaborators' findings by discovering, linking, and visually representing the common work. This approach is based on research in social interaction such as Clark and Brennan's [17], that showed the importance of shared understanding for more effective collaboration. Subtle visual cues enable analysts to gain awareness of each other's findings, hypotheses, and evidence with minimal disruption [34, 45]. This is what I called *partial merging*. Then, if an analyst wants to more closely monitor others' work the *full merging* option is available to integrate others' findings directly into his/her workspace. Partial and full merging are described in detail below. When fully merging views, the layout computation only updates positions of common nodes (those in common between the user's graph and their collaborator's). Common nodes are overlaid and their edge shapes are recomputed as necessary. All remaining nodes (i.e. uncommon nodes) are placed where they originally were, and the user can move them if necessary. I avoid making automatic changes in the layout of the local graph in order to preserve the user's mental map of nodes and relationships.

As a proof of concept for this study, I used Jigsaw to extract multiple lists of entities (People, Locations, Organizations, Chemicals, Events, etc) from the document corpus as well as the list of evidence items to ensure that identical entities and evidence items could be found and merged. There are many ways to improve both the visual representation as well as the merging algorithm used for LCW. However, the main goal at this stage of this work was to demonstrate the value of the LCW concept. In Chapter 6 I further discuss how to improve and extend the LCW technique. There are algorithms to merge entities that are named differently but semantically related (e.g., lexicon chains to find synonyms and related words).

5.3 System Design

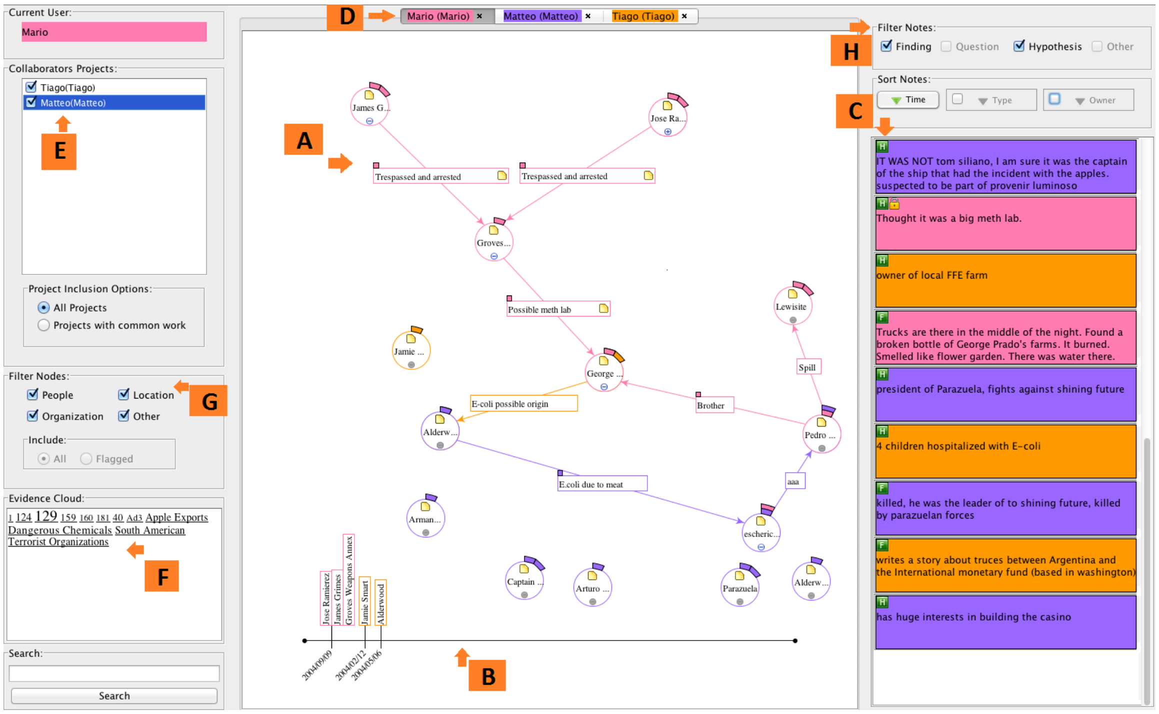

Figure 5.1: Screenshot of CLIP. A) Graph pane, to create a network diagram of people, locations, and events, B) Timeline, to see the timeline of events, C) List of notes, easy review of all notes in one location, D) Tabs, to see collaborators' views, E) Merge option, to choose a collaborator's work to be merged with your own, F) Evidence cloud, to see the list of evidence and their frequencies, G, H) Filtering and sorting options.

CLIP, a Java-based prototype, was designed to facilitate collaborative analysis by providing a visual space for team members to record and share their hypotheses, conjectures, and evidence. CLIP's design takes into account the various design guidelines outlined in Chapter 2. For example, CLIP enables analysts to record externalizations in structured formats (graph, timeline, and notes, see Chapter 2 section 2.4.1), addressing the need for externalization and schematizing. CLIP also enables analysts to work individually but merge their findings (see Chapter 2 sections 2.5.3 and 2.5.4), supporting the principle of shared and individual workspaces. CLIP facilitates validation of results by enabling people to add evidence to each finding and by visually representing evidence in the evidence cloud. Most importantly, I implemented LCW to enhance awareness by revealing relationships between collaborators' externalizations. Collectively, these features help analysts to see who did what and to follow the trail of how other analysts came to a particular conclusion.

CLIP's design is illustrated in Figure 5.1. To facilitate same-time collaboration, CLIP supports awareness through user-controlled sharing of work amongst team members, with colour coding to indicate who did what. The specific scenario for which I designed CLIP is to support team-based analysis of a document collection for solving a mystery task. However, I envision that the design ideas in CLIP could be applied to other collaborative analysis scenarios. Rather than directly visualizing the document collection, CLIP visualizes and links the team members' externalizations relevant to their analysis. Interesting entities (in this case, people, places, or events described in the various documents) can be externalized as nodes, and relationships between events as links. Each recorded entity can be optionally linked to free form text notes and to a timeline. To facilitate same-time collaboration, CLIP supports awareness through user-controlled sharing of work amongst team members, with colour coding to indicate who did what.

In this study, participants were asked to take the role of intelligence analysts solving a mystery task from the VAST 2006 challenge. In the following sections, I describe a scenario related to my study task, illustrating CLIP's use. This is followed by a description of CLIP's features related to externalization and awareness support.

5.3.1 Scenario

Laura, Alex and Mary are reviewing a large set of documents to solve a mystery task. Laura has been focusing on a suspicious event at the 'Silo'. From the article she finds that 'George Prado' is up to some illegal activity, maybe running a meth lab. She suspects 'George Prado' is the key person, so she records him as the main suspect and creates a node with his name. In addition, she creates a note containing her hypothesis that "George is probably running a meth lab". This could also bring George Prado to her collaborators' attention. Afterwards, she starts gathering more evidence to support or refute her hypothesis about him. Later, she finds an interesting article about the 'Silo' and she creates a node for 'Silo'. As soon as the node is created, the visual glyphs on the node inform her that that Mary has also been investigating the 'Silo' event and has recorded information about it. From there, Laura gets interested to see what else Mary has found so far. She opts to merge Mary's entire graph into her own view. She discovers that Mary has found interesting relationships involving the 'Silo' event. Tracing Mary's work, she finds out that both Alex and Mary have collected substantial evidence of the terrorist group's links to George Prado.

Laura decides to also view Alex's full work by merging it with her own. Looking at notes made by Alex, she realizes that 'George Prado' is up to something bigger than running a meth lab. She opens a discussion about 'George Prado' running the 'FFE farm' and from there, they connect the facts and validate their hypotheses. At the same time, Mary records a new finding that shows George's brother is a security guard at the 'Annex', a chemicals warehouse. As soon as Laura and Alex see this new finding in their own views, they start talking to Mary and sharing all their findings so far; in this discussion they realize that the Prado family are probably supplying chemicals to the terrorists.

5.3.2 Externalization

The importance and benefits of recording externalizations (notes, findings, etc.) has been emphasized by many researchers (e.g. [30,47]). My first observational study (See Chapter 3) emphasized the importance of note taking in collocated collaborative VA.

CLIP provides space for recording and visually representing important entities and relationships, the time order of events, and free form text notes. Each recorded entity takes the form of a node in a node-link graph (Figure 5.1 A). Each item is indicated by a unique colour corresponding to the owner. Initially, each user logs into their instance of CLIP by selecting a username and a colour. Evidence can be attached to each node. According to previous research [55], returning to original sources and checking the references can be very tedious and time consuming. Attaching evidence to nodes in CLIP helps analysts easily return to original sources to verify accuracy of reported findings. Each node represents an entity and has six main attributes: text, type, note, image, date, and evidence list. Only the text and evidence list (at least one evidence document) are required for node creation; other node attributes are optional and can be updated any time.

To add a new node to the graph, a user enters information in a popup dialog (Figure 5.2). Each list in the dialog is pre-populated by the values that were extracted from the document corpus. To enter values that do not belong to any of these categories, there is a text box, which is labeled as "Other". To assist users with data entry as well as improve the discovery of common work, I implemented an auto-fill feature for this text box that provides suggestions (name values such as chemicals, events, etc.). A date stamp (null by default) can be attached to the node (Figure 5.2

B). Any node with a non-null date stamp will automatically appear on the timeline. A date stamp can be attached to any node type (i.e. Person, Location, Organization, or Other). The rationale behind this design was to enable users to associate entities of all types with time if needed. Interestingly, in this study I observed that many participants associated people and location names with dates, perhaps as a shorthand way to represent an event. To attach evidence to a node, a user can select documents from the evidence list (Figure 5.2 C). The current design permits attaching up to twelve documents to a node as evidence. The dialog also contains a text area for recording a note (Figure 5.2 D). Note content type (Hypothesis/ Question/ Finding/ Other) and scope (Public/ Private) can be set as well. To assist a user to easily identify private versus public notes, a lock icon appears on the private notes. Finally, an image can be attached to a new node(Figure 5.2E).

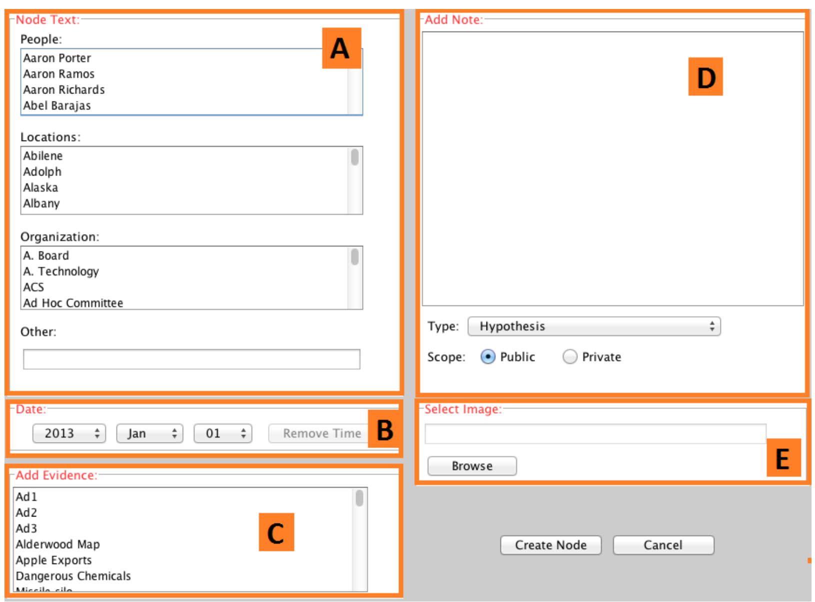

Figure 5.2: Dialog for creating a new node. A) Select node's text, B) Add date, C) Add evidence, D) Free form note, E) Add image.

Each graph node has a toggle button (+ or -) that controls the visibility of the node's children (if any). Collapsing a node collapses all the branches that stem from the node. This feature makes it possible to collapse/ expand the graph from any given node, improving scalability. Users can filter nodes based on the node's type property (people, location, organization, and other) via checkbox widgets (Figure 5.1G). This enables them to both hide parts of graph if required (improving the scalability) and rapidly filter view to a subset of specific types.

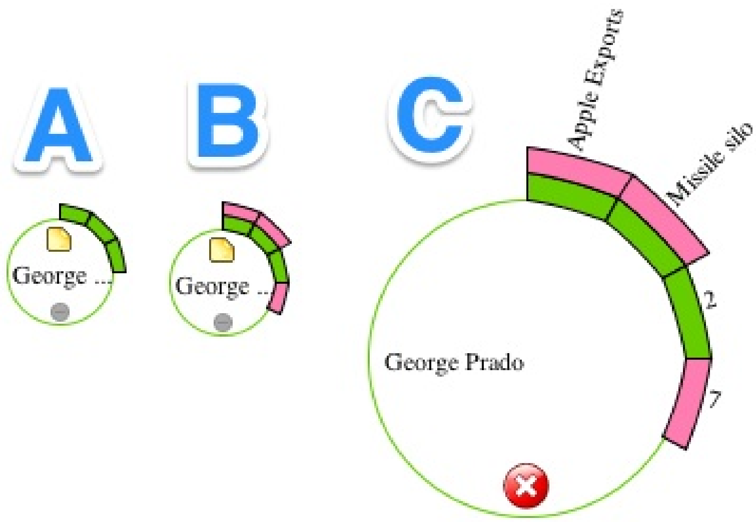

Figure 5.3: Node details. A) View of a node before a collaborator creates the same node and B) after a collaborator creates the same node. Visual cues (colour coded segments) indicate it is a common node and reveal the common and different evidence. C) Enlarged node to see the details. Enlarging a node also highlights related items in other views, such as notes in the note list, timeline items, and evidence.

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 5.3 shows a node's design. Node text is placed in the middle of the circle. A yellow note icon above the text indicates that a note is attached to the node (Figure 5.3A). If the evidence list is not empty, there are segments drawn on the outside of the node circle, one for each evidence document attached to the node. These visual cues provide a quick overview of each node. Segments in a collaborator's colour (Figure 5.3B and 5.3C) represent evidence found by a collaborator, one of the ways we reveal LCW.

By default, all notes are placed in the note panel (Figure5.1C). Each note can be closed by the user later. When closed, the yellow note icon (inside the related graph node) changes to red (to be easily recognized in such a small size) as a visual indication. Re-opening the note reverts the colour to yellow. Users can sort and filter notes by note type, time, and owner (Figure 5.1H). When a node with a Date-stamp attribute is created, the system automatically places a box with the node's text on the timeline (Figure5.1B). The timeline represents the temporal relationships between entities. Items on the timeline are ordered chronologically from left to right and items with similar date stamps are grouped. Inspired by previous suggestions [15, 47], this feature enables analysts to record and review time-dependent findings.

The evidence cloud (Figure 5.1F) is a tag cloud of the documents attached to nodes as evidence. Font size indicates the frequency with which documents have been attached as evidence. Content of this view is based on information included in the workspace.

If a user opts to include all collaborators' graphs, the evidence cloud includes all evidence across the group. This view enables users to easily identify documents that had been noted as relevant, and to reveal something about document importance based on the frequency. This view was added to the prototype based on participants' feedback from a pilot evaluation.

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. Clicking on a document name in the evidence cloud highlights nodes that contain the selected document as attached evidence.

5.3.3 Awareness Support

To support awareness of collaborators' activities, instances of CLIP that are running on different machines communicate in real-time to share information. Distinctive colours are used to distinguish work by different people, similar to Cambiera [45].

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 decorated around the node (colour coded segments). Evidence lists of the local node and the collaborator's node are combined, and repeating evidence segments are stacked up. Figure 5.3 shows a node 'George Prado' before and after a collaborator adds evidence. The colour of the local node is green and the colour of the collaborator is pink. Figure 5.3 A shows the node before partial merging. Figure 5.3 B depicts the same node after partial merging. In this case, the collaborators have two evidence items in common and other evidence items identified by only one person or the other. In addition to these visual cues, CLIP automatically combines all the collaborators' notes related to the node (ordered chronologically by default). By right clicking on a node, users can enlarge it (Figure 5.3C) to reveal more detail. Enlarging a node automatically highlights all related items (i.e., timeline items, notes, and evidence).

Tabs: Each tab in CLIP (Figure 5.1D) encompasses a view of the analysis work in progress 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 5.1 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 5.1E 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.

5.3.4 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. Following one of these events, a message is sent to peers informing them about the latest changes. The receiving instances of CLIP update their views accordingly. Colour is assigned to each CLIP user at the time of application launch. Each user also is required to enter a username at launch time. A user's colour and username are included in the message sent.

Upon receiving a message, the receiving end takes two actions: 1) It 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) The collaborator's work is compared against the local work in search of common entities, which are then merged in the display. In this version, I consider two entities related if the node text is the same (e.g. two nodes with text "George Prado").

5.4 User Study

I conducted a user study to gain a better understanding of LCW's effects on collaborative sensemaking. I employed a between-subjects experimental design to compare CLIP to a baseline tool with the LCW features removed. Details about this study including consent form, introduction to the task and system, tasks, and follow up interview questions can be found in Appendix C.

5.4.1 Participants

Forty-eight participants (sixteen groups of three, eight groups per condition) took part in this study. I recruited grad and senior undergrad students from varied disciplines. To simulate work situations and create a comfortable environment, group members were required to know each other and have previous teamwork experience in a school or work project. To mitigate the possible impact of using students, I mainly selected participants who had some experience with data analysis.

Participants' age ranged from 20-60 (average=28). There were 37 grad students and 28 males out of 48. I assigned groups randomly to each condition. BT: Age range 23-54 (average=30), grads 18/24, male 14/24. CLIP: Age range 20-60 (average=28), grads 19/24, male 14/24.

Participants were compensated with $20 each. To encourage active participation, I provided a small financial reward ($30) for the team with the highest score.

5.4.2 Dataset and Scenario

I employed the "Stegosaurus" dataset from the VAST 2006 challenge [114]. This synthetic dataset contains approximately 240 documents, mostly news articles plus a few maps and supporting pieces of information. The documents describe approximately 3000 entities. The scenario involves finding a hidden chemical weapon production. I chose Stegosaurus because it is a standard task to evaluate visual analytics tools, and has been used in other studies [1, 47, 85]. I was also careful to select a task that represents a real life scenario but can be solved by non-experts.

The dataset contains a scenario that reveals the first clue. From there, analysts are challenged to work through the dataset and iteratively search and filter to find the ten most relevant documents. Similar to real life scenarios, there are distractors that could point analysts in the wrong direction. Dataset authors estimated that the plot could be solved in about 2-6 hours with standard tools [114]. However, I ended all of the sessions at 90 minutes as in [47].

Figure 5.4: User study setup and physical arrangement.

5.4.3 Apparatus

The experimental setup included two iMacs and one 17" MacBook Pro, arranged as shown in Figure 5.4. Participants were collocated and therefore could speak to each other and look at each other's screen if they wished. The physical arrangement was determined through pilot studies where I experimented with several different options to find an arrangement that was comfortable for participants. First I tried arranging the group members within a U-shaped table, so they could easily look at each other's screens; however I noticed that they did not have much discussion. Therefore I arranged them around a table with three laptops to simulate most current work practices, but received many complaints about the small display size. This led to the final larger-screen setup. I expected participants to complain about the screens blocking their views in this configuration, but they reported that it was very practical. I compared CLIP against a Baseline tool (BT). All participants in a group used the same version, either CLIP or BT. BT was identical to CLIP, except that we removed the LCW features (i.e., partial and full merging, as defined in section 5.3.3). I emphasize that BT still contained some awareness features; specifically, collaborators could still examine each other's work through tabs. I kept the tabs because they are similar to what many systems provide currently; for instance, Jigsaw's Tablet view [28] allows analysts to take notes in schematic form but offers no collaborative options to share between concurrent users. This approach also allowed me to specifically investigate the effects of the LCW technique.

5.4.4 Procedure

Prior to the user studies I run three pilot studies. This was mainly to ensure that 1) the analysis tasks were clear and understandable, 2) discover possible bugs in the prototype tool (both conditions), 3) the introduction to task and prototype, follow up interview questions and other instructions were adequate and unambiguous, 4) physical arrangement of displays was appropriate.

I assigned groups to conditions at random. The procedure for both versions was the same. I began with a tutorial on the system's features (15 minutes for BT, 20 minutes for CLIP). I asked participants to try out the system's features with a different sample data set. An observer was present to answer their questions and help them to experiment with all the features. Then, participants received background information about the task and started by reading the scenario, which provided the first clue. Documents were all digital, and all participants had access to all documents. Participants used Mac's Spotlight to search the text corpus. To search within a document, they used Microsoft Word's search functionality. They recorded their results into CLIP or BT.

I ended the study whenever the teams were confident and ready to present their results, or at 90 minutes, whichever came first. Then I asked groups to write a report of their findings and hypotheses. Following the task, I conducted an open-ended interview with each group to discuss the system's features, their challenges, and suggestions to improve the system.

5.4.5 Measures and Hypotheses

In this section I summarize measures and hypotheses related to each of my research questions. I gathered data from five different sources: videos, interaction logs, the final written report submitted by each group, screen shots of visual elements created in CLIP or BT, and notes taken by the observer. All the sessions including debriefing and interview were audio and video-recorded. In total, I gathered 96 hours of video that includes the 16 groups' analysis and follow-up interviews. I used Transana [23] to analyze the videos and measure the total conversation time for each group.

Performance

To measure performance (RQ3a), I analyzed groups' written reports. Using the same scoring scheme as Isenberg et al. [47], groups received positive points for facts they had connected (maximum of 11) and negative points for wrong hypotheses. 11 was the maximum possible score (i.e., all the facts were successfully discovered and connected) and a negative score means that the group uncovered few facts and produced incorrect hypotheses. In addition, and similar to [47], I also counted the number of discovered relevant documents as an indicator of performance. Successful completion of the task was partly related to participants' ability to find the 10 most relevant documents in the corpus and connect the facts within them. I analyzed the screen shots and logs to obtain the number of relevant documents discovered by each group. I hypothesized that CLIP groups would have better results on performance measures, as follows:

- H1: CLIP groups will have better performance than BT: higher task score and greater number of relevant documents found.

Communication, Coordination, and Awareness

I transcribed all the conversations to quantitatively measure communication effectiveness (RQ3b), coordination (RQ3c), and awareness (RQ3d). Using an iteratively built coding scheme, I categorized each instance of conversation. I define an instance of conversation as one or more consecutive statements by a single individual. I chose to code instances of conversation because other possible units, such as sentences, are difficult to clearly delineate in oral conversation. The coding scheme was comprised of seven different categories (DH, RV, CO, SA, VF, QF, and RU). Table 5.1 depicts each code, along with its definition and example. Conversations were coded as DH whenever group members were engaged in a discussion trying to connect the facts and generate hypotheses. This was different from VF (verbalizing findings) when they were not actually connecting facts, they were only stating findings that they found interesting. This usually involved reading parts of a document out loud or reporting a summary of a finding. Referring to the visualization tool (RV) represented instances where participants orally referenced visualization elements such as nodes or notes. The seventh code, RU (Relevant but otherwise uncategorized), was used for any instance of conversation that was related to the case but did not fit within any of the former six codes. Sometimes a single instance reflected more than one code. For example there were instances when participants were referring to the visualization and then they started to have a discussion about their findings and tried to connect them together. I coded these instances as both RV and DH. Other instances of double coding included RV and CO. Therefore, counts of the codes are not mutually exclusive. Over 2800 instances of conversation were coded using the scheme. I did not code conversations between group members and the experimenter.

Code	**Description**	**Example**
DH	Having discussion or generating hypotheses	"So is this a good conclusion to draw from what we have come to so far. US government is supplying the rebels with Lewisite."
RV	Referring to the visualization tool	"Link that with your apples. I will make a new node, linking Parazuelan."
CO	Coordinating the group	"Lets divide the work now, I will search for apples you look for flowers."
SA	Seeking awareness	"What do you guys got?"
VF	Verbalizing findings	"Former farm worker Francisco Dorado formed Shining Future in 1988."
QF	Questions about findings of another group member	"What did you find about apple bursting?"
RU	Relevant but otherwise uncategorized	"Oh okay found that article."

Table 5.1: Communication coding scheme.

Two independent coders coded the conversation data. I assigned groups randomly to each coder. Each coder coded 10 groups (5 CLIP and 5 BT groups) with 4 overlapping ones. Inter-coder reliability was 0.91, calculated using Krippendorff's alpha.

Research has shown that fully sharing the work across the group can trigger discussions that are focused on solving the problem (e.g. [3]). Referring to the visualization also can enhance communication [36]. Because LCW should enable collaborators to more easily integrate their findings and discuss a shared view of their externalizations, I expected that CLIP groups would discuss more facts and hypotheses (DH) and refer to the visualization more often (RV):

- H2: CLIP groups will have more instances of DH and RV than BT groups.

I coded coordination (CO) utterances as those where collaborators tried to coordinate the group activities by dividing the task, documents, the search (e.g., "You search for flowers and I will search for apples"), etc. According to prior research [26], I expected CLIP groups to better coordinate their work. I argue that if the tool supports better awareness, collaborators will be able to coordinate their actions at a much lower level of granularity. That is, instead of simply doing a high-level division of work at the beginning and then sharing findings at the end, collaborators will be able to continually adjust their task division as the work progresses. Therefore, I expected to see more CO instances with CLIP than with baseline:

- H3: CLIP groups will have more instances of CO than BT groups (because they will coordinate at a lower level of granularity).

In order to measure awareness, I coded conversations that were basically for seeking or sharing awareness about each other's activities and findings. For example, questions such as "Are you guys going forward?" or "What have you found so far?" were coded as seeking awareness (SA). Questions about another group member's finding(s) were coded as (QF), and verbalizing one's own findings as a way of sharing was coded as (VF). The rational behind this coding was that I noticed baseline groups spent more time interrupting other members to ask questions about findings or activities, and more time announcing their findings out loud. These questions and verbalizations could be easily eliminated if they could see each others' findings at a glance (the way they could see everyone's results in a merged view in CLIP).

- H4: CLIP groups will have fewer instances of SA, VF, and QF than BT groups (because they will be less reliant on the verbal channel for awareness).

To further explore awareness, I analyzed responses to the interview question about the extent to which participants were aware of each others' work. I also considered checkpoints, when in the middle of the session I stopped them and asked each individual to explain their findings and hypotheses. Then I asked them whether findings of one group member were surprising to others.

5.5 Results

In this section, I present the quantitative results followed by qualitative findings of the study.

5.5.1 Quantitative Findings

Table 5.2 presents scores achieved by CLIP and baseline groups, as well as results of the communication analysis. It reports the number of instances of discussion of hypotheses (DH), referring to the visualization (RV), coordination (CO), seeking awareness (SA), verbalizing findings (VF) and asking questions about another group member's findings (QF).

Task Performance

CLIP groups achieved considerably better scores than baseline groups, strongly supporting H1. As shown in Table 5.2, scores of CLIP groups ranged from 5 to 11 (Avg=8.25, SD=2), whereas baseline groups were from -2 to 7 (Avg=2.75, SD=2.8). The maximum possible score was 11. A two-tailed t test showed a statistically significant difference between the average performance of CLIP and BT groups ($p < 0.001$).

With the exception of group 3, all CLIP groups achieved 7 or higher. I believe the subpar performance of group 3 resulted from their strategy: they spent considerable time organizing the data chronologically before engaging in analysis.

With only one exception (G3, found 9 out of 10) all CLIP groups successfully found the 10 most relevant documents. On the other hand, many baseline groups could not find all the relevant documents. Only two baseline groups were able to find all of the relevant documents (G9 and G16). Even the top three ranked BT groups (9, 13 and 16) who found 10, 9 and 10 relevant documents respectively were not able to connect all the facts. A two-tailed t test showed a statistically significant difference between the average number of relevant documents found by CLIP (Avg=11, SD=1.5) and BT groups (Avg=6, SD=3) ($p < 0.001$).

Task time was not an important factor. I found no correlation between scores and time ($r^2 = 0.028$), and no difference in average time between the conditions (CLIP Avg= 87.6 min, SD=88, BT Avg=86.8 min, SD=87), probably because the task was quite long and most groups used up nearly all the available time.

Communication

H2 predicted that CLIP would foster discussion of facts and hypotheses (more DH). The results strongly support this hypothesis (see Table 5.2). A two-tailed t test showed a significant difference in the number of DH utterances between CLIP and BT groups (CLIP Avg=116, SD= 28, BT Avg=43, SD=45, $p < 0.001$). Although there was no significant difference in the overall talking time between conditions, the difference in DH means that CLIP groups had significantly more discussions about hypotheses and connections between facts.

H2 also predicted that CLIP groups would refer to the visualization more often, and study results also confirmed this prediction. CLIP groups extensively referred to the visualization tool (RV), significantly more often than BT groups (CLIP Avg=30, SD=55, BT Avg= 7, SD=4, $p < 0.001$).

I also observed that there was more discussion triggered by the system in CLIP groups. This was mostly when participants realized their teammate had done some related work. CLIP groups also had fewer awareness seeking conversations, as described in section 5.5.1

Coordination

H3 predicted more instances of CO in CLIP groups than in BT groups (reflecting more detailed task division). I found a significant difference in the number of CO instances (CLIP Avg=15, SD=16, BT Avg= 8, SD=3, $p < 0.01$). In relation to this, I also noticed many instances where CLIP groups coordinated their work via the tool. To further analyze the effect of visualization tool on coordination, I looked into RV examples that were double coded with CO. For instance, I coded this as RV and CO: "Link my node with your apples and I will make a new node to link Parazuela". This is an example of coordination where collaborators deliberately connected their results through the tool in order to solve the problem. I observed and recorded many of these instances for CLIP groups.

Awareness

One of my hypotheses (H4) was that using LCW would help collaborators to maintain awareness of each other's work with less reliance on verbal communication. Conversation analysis strongly supported hypothesis (H4) that the LCW features in CLIP would reduce reliance on verbal awareness mechanisms. CLIP groups had considerably less awareness seeking utterances (SA) (CLIP Avg=2, SD=1, BT Avg=7, SD=3, $p < 0.001$). It was much easier to figure out who was doing what for CLIP groups by looking at the merged view. There was a marginally significant difference in the number of QF (CLIP Avg=4, SD=2, BT Avg=12, SD=8, $p < 0.06$). CLIP groups also verbalized their findings significantly less than BT groups (CLIP Avg=5, SD=

Ver.	Group	Score	DH	RV	CO	SA	VF	QF
CLIP	12	11	185	178	57	0	15	5
	5	10	127	76	15	0	10	7
	8	10	124	26	23	1	1	6
	6	8	131	37	16	2	11	6
	15	8	123	15	10	4	4	3
	1	7	116	20	10	2	2	5
	11	7	102	20	20	1	6	4
	3	5	88	65	11	1	7	2
Avg	-	**8**	**116**	**30**	**15**	**2**	**5**	**4**
Baseline	9	7	116	17	10	9	38	27
	13	6	19	5	5	8	19	9
	16	5	114	6	14	7	10	5
	10	2	23	5	8	7	18	13
	4	2	20	5	9	11	21	15
	14	2	13	8	6	9	14	16
	2	0	11	4	5	4	5	2
	7	-2	25	9	4	1	3	5
Avg	-	**3**	**43**	**7**	**8**	**7**	**16**	**12**

Table 5.2: Comparison of performance, communication and coordination of CLIP versus Baseline groups.

5, BT Avg=16, SD=11, p <0.04).

5.5.2 Qualitative Findings and Usage Statistics

Three primary awareness channels were available to participants: oral communication, LCW (CLIP only), and tabs. To complement and elaborate on the quantitative conversation analysis, in the following sections I report qualitative observations and the results of the post-task interviews for each awareness channel. In the interviews, all CLIP users reported being aware of their collaborators' work most of the time.

They all attributed this to use of LCW features, especially full merging. They found partial merging cues to be an interesting notification of common work that helped them to understand who else had related results and evidence. However, all of the CLIP participants attributed their awareness to full merging. Two CLIP groups (G6, G1) indicated that showing collaborators' notes was another important feature that helped them to maintain awareness of each others' work.

In contrast, many baseline groups mentioned that they were not aware of each others' work. Five out of eight groups reported oral communication as their main awareness mechanism. The rest reported that their awareness channels were oral communication as well as using tabs. These results are consistent with our RV, SA and QF findings. Table 5.3 shows the usage statistics of main features.

Stats	**Node**	**Note**	**Link**	**Timeline**	**Tab**	**E. Cloud**
AVG	20/10	22/14	12/7	10/6	52/71	15/11
SD	6/7	6/8	5/8	3/4	50/40	5/12

Table 5.3: Usage statistics of system features for CLIP/ Baseline Groups

Oral Communication

CLIP groups' oral communication focused heavily on discussion of hypotheses, coordination and referring to the visualization. By contrast, for baseline groups, oral communication was the dominant awareness channel, and without it, participants were not aware of each others' activities most of the time. For instance, one member of group 10 asked, "What are you guys doing, why you don't talk?" and a member of group 4 stated, "I felt I did not know what [Participant A] was doing [she was silent most of the time]". A similar result was reported by Wallace et al. [110], who did not provide any form of thinking or note space for participants in their study.

I observed two key problems associated with communicating only through the verbal channel. Sometimes sharing out loud disturbed others. One baseline participant asked her teammate twice to be quiet. Her teammate was trying to share his findings frequently and make sure that they were all aware of each others' work, but she wanted to focus. Instead she decided to write comprehensive notes to share with the others. Another participant stated, "I couldn't read my stuff when others were telling me what they read, it was grabbing my focus".

The second problem with verbal sharing was that if the information was not recorded, it could be easily forgotten. This supports the importance of providing a means to externalize results and to see others' results through merging and LCW. Although speech was the fastest way to get updated about others' work, there were a few instances where key facts were shared verbally, but later on, the group did not report those key facts in their debriefing (and thus received a lower score than they might have). I noticed this in particular for two CLIP groups (groups 1 and 6); session logs showed that those key facts were never entered into CLIP.

LCW (CLIP only)

Participants reported merging individual work (all CLIP groups, 22 group members), LCW (6 out of 8 groups), and LCW of notes (6 out of 8 groups) as the three most useful features of the system. CLIP participants made extensive use of the LCW features, especially the ability to merge everyone's node-link graphs together. This could cut some unnecessary communication, reduce redundancy, and let team members focus on the task better. According to one CLIP participant, "[Merging] made it faster because I knew what everyone was looking at. I could go on the same direction or do something else...it helped us to collaborate more closely, when you're not paying attention that much, especially I said, 'Hey, what about this?' and someone else is like, 'We've already done this,' and you can just look at their graph. And I can connect my stuff to theirs".

Looking at the most successful groups (12, 5 and 8, all using CLIP), I observed that the strategies they had in common were a clear work division and extensive use of CLIP's LCW features, suggesting that these are good predictors of success. Even though these groups had different leadership styles, they constantly divided the workload among members. According to the system log, these groups also made intense use of CLIP's features to coordinate their activities. Participants systematically merged their partners' work into their own view and linked their work together.

Participants reported that the merged view helped them find important results through others' work. It also helped them figure out what keywords to look for, and inspired confidence. Participants said, "I noticed that some of my most powerful points... he also had them. I could see the two colours on it. That gave me confidence," and, "For common nodes, I was looking at the evidence. If they were different from mine, I was checking them as well. Common items made me confident and helped me to keep going." CLIP groups became quite dependent on the shared node-link graph. For instance, in group 5, B was sharing her findings with A, and A said "Put it down, create a node". Later when B complained that the team ignored one of her findings ("I found it before and I told you!"), A explained why it had gone unnoticed, "Because you did not connect it to my node, so I did not look at it!".

I was curious to see how participants in CLIP groups would choose to use merging. Would they leave the default setting (partial merging) to keep their workspace uncluttered and avoid the disruption of constant updates from other participants' changes? Or would they choose to see everything? Answer: the latter. Most participants chose to set merging on from the beginning and kept it visible until the end.

It was interesting that participants reported that oral communication was disruptive, but CLIP updating the shared view was not. Instead, participants reported merging to be useful for collaboratively exploring the task, sharing important evidence, exchanging documents, and reducing redundant work. During the interview participants emphasized that merging was one of the most useful features of CLIP.

Five groups (eleven participants) in the baseline condition actually requested a merging feature that would put everyone's information in one view. For example, participants stated, "I was not able to make a link to someone else's work, so I could not make a connection," and, "It is hard to remember what the others have registered by checking the tabs, so we would like to be able to draw links between nodes created by different people. It is also good to avoid redundancy."

Only one participant reported a potential negative side of merging, stating, "It was interesting, but it was a double-edged sword. It could help me or push me [in a] correct or incorrect direction."

Tabs and Notes

In addition to oral communication, most baseline groups also relied heavily on tabs for awareness. Some groups, however, used tabs only to quickly check what the others were working on; for example, "I only looked at their tabs when I was trying to find something that they have read. I just wanted to refer to their work but not for everything". Interestingly, tabs were also sometimes useful to CLIP groups, even though they also used the merging feature. One CLIP participant explained that she used tabs to see how other group members arranged their nodes (because CLIP's merge feature recomputed the layout separately for each individual to maintain their mental map).

Notes were valued by both CLIP and baseline groups. Collaborators' notes were accessible via tabs in both tools, and via merging in CLIP. Participants stated that notes provided an overview, enabled them to remember why they had created graph nodes, and allowed them to copy important information from the documents. Several people reported that notes helped them to identify interesting information belonging to others. For example, participants in group 12 stated, "[C:] The notes on the side. I got mostly info from them, to be honest. I would read the notes and go 'Wow, that's cool!' [B:] Yeah, other people were highlighting things that you should read." Similarly, another participant said, "When someone didn't write a good note, I didn't look at what they were doing".

5.6 Discussion

LCW clearly supported groups in this collaborative sensemaking task. CLIP groups achieved significantly better scores (H1), coordinated and communicated more effectively (H2 and H3), and relied on LCW to maintain awareness of each others' work (H4). CLIP groups had significantly more discussion about hypotheses and evidence (DH) and were able to focus their oral communication on discussing the case and coordination activities rather than using oral communication as the main awareness channel. This research extends earlier work on LCW [45], by establishing its value in the sensemaking loop phase of collaborative analysis, not just the information foraging loop, and demonstrating how it can be applied to externalizations. CLIP also illustrates how the LCW concept can be employed within a collaborative thinking space.

To better explain the effects of using CLIP on teams' collaboration, I derived a collaboration model for CLIP groups based on this study's results (see Figure 5.5). Similar to the model in [26], this model shows how awareness plays a critical role that enhances communication and coordination activities in collaboration. From Figure 5.5, we can see that recording externalizations and automatically sharing those externalizations via LCW resulted in increasing awareness. Increased awareness in turn enabled groups to coordinate their work at a deeper level. Being able to see others' results triggered discussion and this assisted teams to better formulate their new questions and hypotheses (QH Formulation). There are mutual effects between awareness, coordination and discussion which means each one will influence the other.

QH Formulation and coordination of activities initiate and direct new investigation.

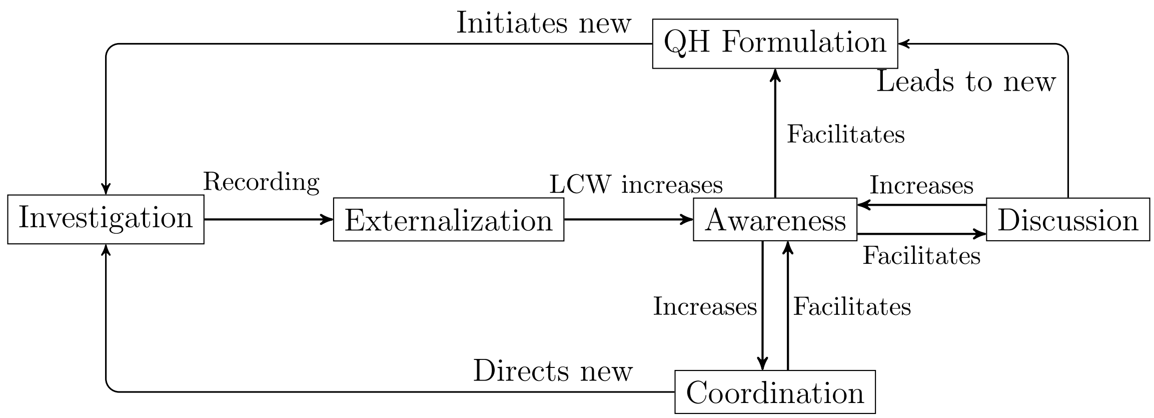

Figure 5.5: A representation of the dynamics of CLIP groups' collaboration. This model emphasizes the role of LCW in increasing awareness and discussion among team members. Awareness leads to better coordination of activities and formulation of new questions/hypotheses, which in turn initiates and directs new investigation.

With the Baseline tool, LCW was missing. In terms of my model in Figure 5.5, this means that the link between Externalization and Awareness was effectively broken. Collaborators did still maintain some awareness of each other's findings and activities through oral communication, but this mechanism was less effective. Reduced awareness in turn had detrimental effects on the teams' coordination, discussion, and investigative activities.

While CLIP provided an effective thinking space for the intelligence analysis task in this 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 explicit support for information foraging. Combining CLIP with a complementary tool like Cambiera [45] may be an effective way to support both phases. I strongly suspect information foraging support would be crucial when dealing with a larger document set. Awareness cues related to information foraging could present details about each collaborator's current activities. For example, CLIP does not show whether a specific collaborator is currently reading a document or entering a search query. Another limitation of CLIP is that it does not automatically identify entities or relationships between documents; I extracted and encoded these manually for the purpose of this study. Many algorithms exist for finding these relationships automatically. For this reason, I would like to integrate CLIP's thinking space and LCW features into a document analytics tool such as Jigsaw [100] that automatically extracts entities and relationships from the documents. Integrating CLIP with tools like Cambiera and Jigsaw has potential to further enhance awareness, resulting in more effective communication and coordination. However, future research will need to investigate whether it is possible (and how) to provide this extra information without causing information overload.

Another way to extend CLIP would be to add visual indications that distinguish past work from current changes. This was recommended by one of the groups. Although the dynamics of the node-link graph show the evolution of the team's findings, it is not clear at any given point in time which are the most current changes. Adding a list of new changes, along with use of visual cues such as colour saturation to indicate a node's age, could help analysts to easily see the new changes at a glance. Although a timeline of actions could provide the same information, I argue that integrating this information directly into the workspace should be easier for analysts to follow as there is less cognitive disruption from solving the task. Another interesting feature suggested by participants was to create a summary evidence file to help with publication of the results.

Scalability is another important issue. In the current design of CLIP, collapsing a node collapses all the branches that stem from the node, improving overall scalability and providing flexibility. To scale this design to very large and complex problems or datasets other than document collections, different visual representations might be needed. The visual structures (e.g., node-link graph) 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. I predict that the 'share-everything' strategy that was successful for CLIP groups in our study might break down at a larger scale. A variation of Branch-Explore-Merge [71] might reduce the number of visual updates since they would only appear upon merging. For small thinking spaces, this may be a significant disadvantage since awareness notifications would be delayed. However, in large thinking spaces, providing awareness notifications in such chunks may cause less visual distraction and reduce the likelihood of small updates being missed. There are also other issues of scalability in the current design. First, while the colour coding works well for small groups (our target), it should be reconsidered for larger groups. Also, for the specific task used in this study, decorating evidence around a node was enough. However, the design might need to change for larger datasets with more evidence items. One possible way to improve scalability could be to encode the quantity of evidence related to a node as the node size. Similarly, the size of notes could adapt to their length.

One interesting question that arises with visual thinking spaces is the potential that they may lead to *groupthink*, a situation where the group fails to consider possible explanations because they too quickly follow one avenue of investigation. It is possible that sharing findings through LCW may discourage a healthy level of independent analysis. I do not have a good way to assess the level of groupthink in this study, in part because avenues of independent thought are neither easy to categorize nor measure. One possible approach to avoid groupthink could be to design tools that promote discussion of alternative hypotheses. Mechanisms to find and highlight disagreements in the findings might be good for initiating such discussions. Research into causes of groupthink and mechanisms to prevent it are an important area of future research. Nonetheless, the much stronger performance of CLIP groups in our study indicates that the awareness benefits of LCW outweigh costs such as groupthink.

Future work should also examine the value of LCW of thinking spaces in a field setting with professional analysts. The study participants were necessarily students, because it is extremely difficult to find enough professional analysts for a lab study. I took care to recruit participants with some relevant data analysis experience (mostly graduate students who were familiar with research), and chose a task that did not require domain-specific knowledge. Nonetheless, student behaviour will undoubtedly differ from that of experts. For example, professional analysts may have established coordination strategies that they could employ and might therefore be less reliant on tool support to assist with coordination. Furthermore, I would like to explore how LCW influences collaborative dynamics over a longer analysis period than a single work session.

Another interesting future direction is to understand how CLIP could be used on a shared screen (e.g., a wall or tabletop). I examined the value of LCW for collocated work; however awareness cues offered by LCW might have even greater value for scenarios where collaborators are distributed in time and / or location. Maintaining awareness of collaborators' activities is generally more challenging in distributed scenarios because of the limited communication channels available to collaborators. LCW could play a critical awareness role in such situations, but this will need to be tested in future studies. It is quite possible that additional features will be needed (e.g., a more extensive note feature that enables threaded discussions) when verbal and / or non-verbal awareness communication channels are unavailable.

5.7 Conclusion

In this chapter I addressed the effects of linking collaborators' externalizations to one another on awareness, performance, communication and coordination (RQ3). In order to do this, first I designed and developed CLIP. CLIP demonstrates how the concept of linked common work can be employed within collaborative thinking spaces to support the sensemaking loop during collaborative analytics. CLIP provides an environment for analysts to record, organize, share and connect results. Moreover, CLIP extends earlier thinking spaces by integrating LCW features that reveal relationships between collaborators' externalizations to increase awareness among team members. In this user study, I compared CLIP to a baseline version without LCW features. Results demonstrated that LCW significantly improved analytic outcomes at a collaborative intelligence task. Groups using CLIP were able to communicate and coordinate more effectively. They were able to use oral communication primarily to discuss the task, generate hypotheses, and coordinate their activities at detailed level, rather than employing it for disruptive awareness notifications. Most importantly, LCW enabled collaborators to maintain awareness of each other's activities and findings and link those findings to their own work.

The most important contributions of this chapter are the design and evaluation of CLIP which demonstrated that automatically identifying and linking common work has value for supporting sensemaking in collaborative VA.

Chapter 6

Discussion and Future Work

In this chapter I further discuss contributions of this research and their applicability in other domains. I elaborate on the proposed models and list all the guidelines that I directly derived from the studies. Then I discuss how collaborative VA tools can help teams to succeed. Next I address threats to validity such as the generalizability of my results and their limitations. Finally, I point out research directions derived from this research. This thesis was an initial step to better understand challenges and needs for recording and sharing analytical results during collocated collaborative visual analytics. There are still many open questions that can be considered as future work. I hope the contributions of this thesis motivate more research to further investigate this field.

6.1 Proposed Models for Processes and Externalizations in Collocated Collaborative VA

In the first study, I proposed a characterization of collaborative VA phases and activities based on the study results as described in Chapter 3. In the second user study I observed similar activities that corroborated the structure and the relationships in the proposed model. This suggests the robustness of the model to characterize the iterative process of analysis activities during collaborative VA.

The activities in the third study were not identical because the context was different and users did not use a visual analytics tool to analyze the data. However, the externalization model (described in Chapter 5) that I proposed based on the third study justifies the placement of record-keeping as a critical activity in the first model.

While these two models seem to be separate, they are actually quite related and corroborate each other. The first model highlighted the importance of record-keeping activities during all phases of collaborative VA for a small group of three. Though the context of the third user study was different from the first two, my findings still showed the integral role that record-keeping combined with LCW played in facilitating collaborative activities. The LCW technique allowed team members to gain insights based on each others' externalizations. This model indicated that supporting collaborative recording and sharing of externalizations can significantly improve the analytical outcomes, leading to better communication and coordination. It is notable that this model also derived from a user study with groups of three.

6.2 Design Considerations to Support Recording and Sharing of Analytical Results

In this section I integrate the design guidelines proposed by different phases of this research to support the sensemaking process during collocated collaborative VA for small teams (up to size three). However, some of the proposed guidelines may have wider applicability. Since these implications address different settings in collocated collaborative VA (e.g. tabletops, desktops, etc), they should be considered if they apply. For instance, guideline number 4 suggests improving input mechanisms for tabletop applications.

In Chapter 2, I reviewed existing design guidelines for collaborative VA proposed by Scott et al. [93] as well as the list of guidelines proposed by Isenberg [42]. These guidelines put little emphasis on design for recording and sharing of analytical results. In this thesis, I focus on guidelines that are related to supporting externalization and the sensemaking process. While some of these deign considerations are in line with previous work, they emphasize how to design record-keeping functionalities based on previous guidelines (i.e. the first seven considerations). The most novel design implications derived from my research are the last three: supporting an unobtrusive sharing channel, supporting LCW, and providing visual cues to increase awareness. I expect these to have a broader application for collaborative VA tools.

In this section, I only discuss suggestions and implications that I have directly learned from my studies, rather than those that I conjecture. While I present a blend of higher level and lower level considerations, these design considerations work together to support the higher level task of sensemaking.

(1) Support both independent notes and those linked to visualization states
According to my findings, a record-keeping module should be designed with enough flexibility to support taking both independent notes and those linked to visualization states (e.g. a chart). When the note can be linked to a visualization state, retaining this link and enabling the user to return to the artifact and system state would have a clear benefit. However, there are other types of notes that are not directly relevant to one single visualization state (e.g. reminders). In this case, a notebook style representing all the notes in one single place would be more useful.

(2) Support Schematization, Manipulation and Management of Recorded Materials
Note management and manipulation is also important to consider. Enabling highlighting, grouping, and summarizing of notes will increase their usability [10, 57, 97]. Previous research has shown the importance of providing users with functionality to manage and structure their notes. For example, in the context of education, it has been mentioned that a matrix structure for recording notes is more beneficial than a linear structure [7]. In the first study, I observed that most of the group notes were naturally taken in tabular format (Figure 3.5), which facilitated comparison. In the third study, I discussed the value of searching, filtering and sorting notes and suggested adding abilities to search, sort and filter notes based on different parameters such as time and type (question, hypothesis, finding, reminder, etc.). I also suggested supporting schematization: providing a uniform structure to enable group members to interpret each other's recorded materials.

(3) Support Privacy for Individual Notes and Controlled Sharing
In a collaborative setting, collaborators should be able to control how recorded materials are shared within the group. In my third user study, the design of the note taking module enabled users to mark a note as public or private. It was interesting that even for this collaborative task (as opposed to more competitive task used in the first study), which was also very time critical, users took the time to mark some of their notes as private. In addition the observations in the first study that some notes had a private scope, this suggests that privacy is an important consideration. Providing privacy options might be more critical in some disciplines than in others.

Sometimes records of decisions need to be kept private. For example, in the field of intelligence analysis higher ranking analysts may wish to share findings among the group based on individual clearance. Also in the third study, participants asked for mechanisms for one-to-one sharing when they had specific inquiry intended for only one member of the team. This suggests that users should also be able to share items with a selected subset of the group.

(4) Improve Input Mechanisms for Large Interactive Surfaces

The note taking module for large interactive surfaces need to be improved. In my second user study, I used on-screen keyboards as the input mechanism. However, many participants found them awkward and cumbersome. This suggest that the design for note taking module for large interactive surfaces should be improved. Providing alternative recording mechanisms such as multiple physical keyboards and mice and/ or digital paper could solve this problem in some extent. However, more studies are needed to understated and implement forms of input that closely resemble natural note taking behaviour.

(5) Support Multiple Coordinated Views of Recorded Materials

The design of the system should support multiple coordinated views of recorded material. Users in my studies valued and reported positive feedback about having various views of recorded materials. For example, in CLIP, participants valued the ability to review evidence in relation to artifacts in the node-link graph view, as well as the evidence cloud, which showed the frequency with which evidence documents were attached to nodes. CLIP also provided users with two different views of the recorded entities (graph and timeline views); users reported that it was very useful to relate items across these two views.

(6) Support History Management

In study 2, some participants frequently performed manual save and delete actions on the history. At least one participant reported that having to manually delete unwanted items was much more cumbersome than having to save desired ones. This behaviour indicates the perceived importance of history as a container of important analytical artifacts. The observation of abundant "quick reviews" (i.e. actions on history that did not involve any direct physical interaction with the system) throughout the second study also strengthens this speculation. These results suggest that manually saving items to the history may be a better choice in many circumstances over automatically saving them. Alternatively, an improved algorithm could be devised that would save the analysis state less often and better infer the states that are important to keep.

(7) Support Record-keeping for Different Collaboration Styles

My thesis results suggest that changes in collaboration style should be considered in record-keeping design. Based on my observations during the second and third studies, the design of a record-keeping module should support recording the results while team members are working in both loosely coupled and closely coupled fashions. Furthermore, the design should support an appropriate sharing channel (between either the whole or a subset of the group), allowing team members to share and discuss their results.

(8) Support an Unobtrusive Sharing Channel

Analysts should be able to remotely view (given proper permissions) both current and past work of their collaborators. The sharing channel should provide remote, unobtrusive and non-interruptive access to a collaborator's work. In CLIP, users could view collaborators' work by either using tab views or by merging their externalizations into one view. In each case, the viewer would not interrupt the collaborator's work.

(9) Support Linked Common Work (LCW)

Based on my findings, LCW can noticeably improve team awareness, coordination, communication, and analytic outcomes. Therefore, record-keeping modules should provide appropriate support for automatic discovery, linking, and visual representation of related findings among group members. The underlying implementation of LCW may differ based on the domain, task and data. Yet, the benefits of LCW should transcend these factors and can be applied to other tasks and domains.

(10) Provide Visual Cues to Increase Awareness

In addition to visually representing the common work (LCW), other visual cues can increase awareness among collaborators in VA tasks. Participants in study three asked for a mechanism to distinguish past work from current changes. One possible solution could be to employ colour transparency as a temporal indicator of the work progress of collaborators. In a merged-views situation, nodes of a collaborator's graph could be drawn with different levels of transparency to indicate recency (e.g. most recently added nodes could be solid and older notes could be progressively more transparent).

In the future work section I discuss directions to extend these design guidelines for different display configurations and different collaborative settings.

6.3 How Collaborative VA Tools Can Help Teams to Succeed

The main focus of this thesis was on supporting a major task during the collaborative sensemaking process: recording and sharing the analytical findings. From my results I have made suggestions for the design of future tools. While I have discussed guidelines related to recording and sharing of analytical results, there is another important aspect to consider for designing collaborative VA tools: behaviour that can be encouraged by CSCW technology to help teams to succeed.

In the third study, the significant improvement of performance, communication and coordination among CLIP groups inspired me to think of behaviour that collaborative technology could encourage to help teams to succeed. My suggestions here are based on a combination of my observations in the 3rd study as well as previous literature results. The first suggestion is to encourage closer collaboration, as widely discussed by previous research [8, 47, 55, 109]. Isenberg et al. [47] and Vogt et al. [109] found that teams who collaborated more closely were more successful, and Bradel et al. [8] reported that members in more successful teams understood each other's work better. Groups should not always work in a closely coupled fashion, however. Kang and Stasko [55] found that in a long term project, collaboration was loose during information collection but tight when synthesizing findings and writing a report. 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 to each investigator about their collaborators' activities and findings. Paul and Reddy [82] suggested providing support for both action awareness and activity awareness, showing actions that led to a particular activity. Therefore, mechanisms to increase awareness are needed. The results of the CLIP user study suggest that the LCW technique is one possible solution to increase awareness and encourage closer collaboration. Further research is needed to investigate other advanced mechanisms and their effects on collaboration mechanics.

A second suggestion is to encourage group coordination, by helping team members to better divide their tasks and share their findings. Balakrishnan [2] also highlighted the importance of coordination. She proposed a model of the collaborative sensemaking process that consisted of two main components. The first component was sensemaking (the process of finding meaning from information). The second component was coordination, which steers collaboration toward successful problem solving. According to Balakrishnan, “Successful teamwork depends on effective coordination. Successful coordination involves communication among partners, shared resources, a shared understanding of the group’s goal, [and] an agreed upon overall strategy or plan.” I evaluated this suggestion in the third study and the results indicated that increased awareness, shared resources, and shared findings can significantly improve both communication and coordination. Further research should investigate other mechanisms and techniques to improve coordination among team members.

6.4 Threats to Validity

It is well known that the choice of methods and study settings can directly affect the results and their validity. In section 1.3 I justified the choice of methods and discussed approaches I took to offset possible flaws. Nonetheless, every single choice in designing and evaluating a lab study can influence the results. As expected, results of the studies conducted in this thesis are subject to some caveats. In this section, I describe the limitations of this thesis and address the empirical validity of the results by describing threats to construct validity, internal validity, external validity and reliability [20, 70].

6.4.1 Construct Validity

Construct validity “focuses on whether the theoretical constructs are interpreted and measured correctly” [20]. This threat to validity is mostly applicable to the third study because in this study, I used a set of metrics to measure concepts such as awareness between team members. I developed metrics to measure concepts that were often difficult-to-quantify: awareness, coordination, communication and referring to the visualization tool. Although I paid careful attention in choosing and developing each metric to accurately and effectively measure what I was interested to measure, these metrics are not perfect. For coordination, I was interested to measure coordination but what I coded were instances of discussing coordination. This was only one aspect of the coordination process but at least it was a visible and measurable aspect. To measure awareness, I could not directly measure awareness because it is an internal mental concept. I coded SA as instances where participants were seeking for awareness. Also I relied on their verbalizations, which may not accurately represent the underlying cognitive processes. For discussions among team members, perhaps I captured most of them because it was obvious when they were having discussions. However, the topic of discussions sometimes was hard to capture due to a few inaudible/distorted segments of the videos. RV (referring to visualization) however was one of the harder ones to capture. I might have missed some instances, because to code this I had to simultaneously look at videos, screen captures and systems logs to make sure they were referring to a recorded material.

6.4.2 Internal Validity

Internal validity deals with "the degree to which results of a study permit you to make strong inferences about causal relations" [70]. I aimed to carefully control differences between conditions (e.g. in CLIP study the only differences between conditions were the presence and absence of LCW, otherwise the systems were identical). But, there are still some issues that present possible threats to internal validity.

In the first study two coders coded the data together, however, independent coding may have led to higher reliability. Therefore, learning from the first study, in the second study two observers coded data independently. Similarly, in the third study, two coders coded the conversations independently. We went through a series of iterative codings, in terms of disagreements we discussed the results until we came to a consensus. In addition, I iterated between qualitative and quantitative analysis to ensure that results were supported from findings of both types. Moreover, for the Likert style questionnaires in the second study, I closely followed accepted standards to alleviate acquiescence bias. In the third study, I proposed LCW as a new technique to increase awareness among team members. However, I cannot separate out all of the features of LCW to know which specific ones were helpful because I did not evaluate them separately.

6.4.3 External Validity

External validity "refers to how confident you can be that the findings of your study will hold up upon replication, and how confidently you can predict both the range over which your findings will hold and the limits beyond which they will not hold" [70]. Here I elaborate on issues that deals with the generality of results as well as their practicality and validity.

First of all, in all three studies I involved students as the participants, because it is extremely hard to find real analysts for such lab studies. While in this research I paid careful attention to the study setting (e.g. hiring collaborators with possible joint work experience, choosing appropriate tasks, etc.), it is not clear to what extent the student participants behaved like real analysts. I also chose tasks that were carefully modelled after real domain tasks, involving experts in the design of the task or using standard benchmark tasks. In the first study I focused on business students who have some knowledge of the domain. In addition, I chose datasets that were simple enough for non-professionals but complex enough to challenge people and engage them in a collaborative effort to solve the task. However, domain experts are trained for specific analysis tasks, they usually need to manage higher volumes of information, they likely have access to a greater variety of datasets and tools, they have preexisting relationships with their collaborators, and they know how to handle their responsibilities within their work environment. Moreover, there are many conditions in real life that affect groupwork such as stress, and time pressure, which we did not deal with. These differences all limit the value of students and the value of a lab study as a model situation.

Another important factor in collaborative studies is whether or not team members should have collaborative work experience before the study or not. In real life situations, collaborators build experience working with one another, which gradually improves their communication. To emulate real work situations as closely as possible, in the first study I set a condition for group members to know each other and to have some previous collaborative work experience. This made the recruitment process very difficult and time consuming. In the second study, I decided to recruit pairs who were familiar with visual analytics but did not necessarily know each other beforehand. This made recruitment easier but may have reduced the similarity to real world collaborative situations. Participants in this study showed great interest in the record-keeping module. However, the record-keeping functionality and awareness features did not lead to observable increases in discussion among team members. I suspect that real analysts would have had more discussions around their findings. In the third study, I again recruited people who knew each other beforehand and had some previous joint teamwork experience.

The fact that group members were not always team members through all studies could be a threat to validity and makes it slightly more difficult to make comparisons across the studies. In addition, it should be noted that people who know each other in a personal context (such as friends) have a different kind of relationship than people who know each other as colleagues. Therefore, in future work it will be important to conduct studies with domain experts and with groups of colleagues that have real joint work experience, to ensure the validity of my results.

Another important factor is the group size. Collaborative sensemaking processes are likely to vary with the group size. It is well known that size impacts the social dynamics of a group (e.g. [99]). While most of the prior studies in collocated collaborative VA studied pairs (e.g. [8, 47, 109]), I examined group sizes of three and two. Studying group sizes larger than two was particularly important in the third study because as the size of group gets larger the complexity of sharing and understanding each other's results became more challenging. Choosing groups of three in the third study (solving a complex intelligence task) was very helpful to investigate the LCW technique with a more complex and larger number of findings. This study also corroborated the results of the second study conducted with pairs. Thus, I suspect that results of my thesis can be generalized to other small groups as well. However, these results are only valid for small teams (2-3) and can not be generalized beyond these very small group sizes without further studies.

Although the main focus of the thesis was on understanding the collaborative sensemaking and challenges to support recording and sharing of results in general, the selected application domains might be a limitation to generalizability. In the first study, I chose to focus on the business domain. I suspect that collaborative use of visualization tools will be similar for other group decision making tasks, but it is possible that I observed some peculiarities unique to business. In the second study I evaluated the effects of providing support for record-keeping in business domain. However, I identified nine primary actions which maybe domain-independent and carry some meaningful information regarding users' intentions. While in the third study I changed the domain to intelligence analysis, the results of my study showed the value of support for recording and sharing. The core ideas (supporting record-keeping and sharing of results, and providing an unobtrusive awareness channel) were common in both the second and third studies.

While the design of CLIP might have been specific to document collections used in intelligence analysis, and to the specific Stegasaurus task, the core idea of LCW should be generalizable to other domain problems. The model that I derived from this study (Chapter 5) corroborates previous models in CSCW and highlights the role of supporting awareness for groupware technologies [26]. Yet further studies are needed to determine applicability of these results to other domains.

Another limitation to generalizability of the results is the tools that have been selected or designed for these studies. In the first study, I chose to utilize existing visualization software to study users while interacting with representations of data. However, the users' behaviour may have been influenced by the available technology. In particular, the teams' closely coupled work style may have been influenced by the tool design (owing to the inability to create different charts in parallel) and the tendency of most groups to write notes rather than save charts may have been similarly influenced (owing to the difficulty of saving charts with the tool). For studies two and three, I paid special attention to the design of the prototypes and followed previous guidelines to ensure that these prototypes supported collaboration as well as record-keeping activities. However, the results of these two studies might still have been influenced by the design of the prototype tools. Future user studies are needed to examine design decisions made in these tools and the effects of providing different record-keeping mechanisms and awareness channels on the visual analytics process.

6.4.4 Reliability

Reliability "focuses on whether the study yields the same results if other researchers replicate it" [20]. I provided detailed descriptions for each study to be replicable by other researchers. While the results of coding for studies 2 and 3 had high inter-coder reliability, it is possible that other researchers might find other interesting patterns in the data or code the data differently. Overall, this thesis initiates a research direction to better understand needs and challenges of a small team during the collaborative sensemaking process. While there are certain limitations for each phase of this thesis, the main findings in the last study reflect what I found in earlier studies. In addition, results of each study corroborated many findings reported in previous work. Nonetheless further studies are required to better understand the reliability of these findings.

6.5 Future Work

This research is an initial step towards understanding requirements for recording and sharing results in collocated collaborative VA. However, there remain many promising avenues for future research. All the limitations discussed above call upon further research. In addition, future research can go beyond this list to investigate additional aspects. In this section, I point out some of these opportunities.

6.5.1 Additional User Studies

In section 6.4 I discussed the limitations of this thesis and addressed the empirical validity of the results by describing threats to construct validity, internal validity, external validity and reliability. All of the limitations that I pointed out in section 6.4 could form future research directions. For instance, one possible direction is to repeat all of the studies conducted in this research with real analysts and with teams of experts with joint teamwork experience in the field. Additional studies should investigate larger groups and other application domains. In addition, the form of input to be used for note taking in a shared display situation should be investigated in more detail.

6.5.2 Develop and Evaluate Tools

Another research direction lies in the development and subsequent testing of new visualization tools and features. As Heer and Agrawala [36] point out, there are a number of design considerations that present opportunities for collaborative visual analytics. The next step is to implement and evaluate these design suggestions. I have not implemented and evaluated all of the suggestions that were derived from my studies, but doing so will be useful to better inform the VA community. In addition, there are other recording and sharing mechanisms that have been outside the scope of this thesis and require further research to understand how they will improve collaborative problem solving.

Furthermore, this thesis provides many suggestions on how to support both individual and group record-keeping activities. However, exploring additional mechanisms for record keeping and awareness is still worthwhile and may further improve upon the guidelines.

6.5.3 Extending the Design Guidelines

Other interesting future direction could be to evaluate the usefulness of the proposed design guidelines and explore how they can be extended to other display configurations and other collaborative settings:

Extending the Guidelines for Different Display Configurations

Considering that the three studies used different display configurations, it is important to investigate how the proposed guidelines apply in different display configurations. The proposed guidelines based on the first two studies should be tested for desktops and the ones from the third study should be evaluated for shared displays. Understanding how CLIP could be used on a shared screen (e.g., a wall or tabletop) might lead to new design insights and improved guidelines.

Extending the Guidelines for Different Collaborative Settings

It should not be difficult to think of extending these guidelines to remote synchronous or even asynchronous collaborative settings. For instance, although I examined the value of LCW for collocated work, it might have even greater value for distributed or asynchronous scenarios. Many of the design ideas could applied to these modes of collaboration (e.g. colour coding collaborators' contributions). However, different/ additional communication channels and awareness mechanisms may be needed to support other collaborative settings. Maintaining awareness is generally more challenging in distributed situations because of the limited communication channels available to collaborators. LCW awareness features could play a critical role in such situations, but this will need to be evaluated in future studies. It is quite possible that additional features will be needed (e.g., a more extensive note feature that enables threaded discussions) when verbal and/ or non-verbal awareness communication channels are unavailable.

6.5.4 Enhancement and Extension of the LCW Technique

The CLIP study (Chapter 5) demonstrated that awareness features assisted team members to communicate and coordinate more effectively and achieve better analytical outcomes. This is an interesting finding that addresses one of the grand challenges in designing collaborative VA tools [18]. Therefore, it merits more attention. Future research should investigate how this method could be enhanced and how it could be applied to different domain problems. For instance, what are the best forms of representing linked common work in different domains? How to address scalability when the number of collaborators or the number of findings increases? How should this method be applied to the whole process of sensemaking, including the information foraging loop? Adding abilities for information foraging to CLIP might answer the last question, but there are still many intriguing questions left to pursue for future research.

Another future direction for CLIP could be to notice, differentiate and visualize the contributions of individuals, pairs and the whole group. This derived from my observation of a blend of individual and group effort in terms of contributions to the final solution. A group solution was comprised of different individuals' contributions as well as contributions made by the group or a subset of the group. This was particularly obvious in the this study, because group members had the ability to merge their work. The individual contributions were connected together to form the final solution. The colour coding in CLIP showed who had contributed to the final solution (i.e. it indicated who had created each note). However, it did not visualize and identify findings that had been discussed within a pair or within the whole group. The question that arises here is how we can notice and measure the contributions of pairs or the whole group? This was not the focus of my thesis; however, I speculate that visually representing the contributions of individuals versus the contributions of the whole group or a subset of the group is useful. This will make it easy for experimenters to understand who contributed to the final solution and where. In addition, this information might be useful to the collaborators themselves, and might encourage more equitable participation. In the CLIP study I did not have the baseline case of no sharing. However, it might actually be interesting to see whether the quantity of written contributions, or equality across the groups, would differ between CLIP and a baseline tool that had no sharing mechanism at all. For the above purposes, the current design of CLIP is limited in several ways: it only captures who recorded an item, and does not represent pair or group contributions. CLIP also does not capture discussions or verbal contributions. Future design enhancements are needed to support this functionality. Visually representing the individuals' contributions makes the tracking of past discoveries easier, perhaps it might also encourage contribution. This is an interesting research direction for studies that investigate group behaviour.

6.5.5 Providing Solutions for Groupthink

Another avenue of research left to pursue is investigating groupthink. Groupthink is a situation where the group fails to consider possible explanations because they too quickly follow one avenue of investigation. It is possible that sharing findings through LCW may discourage a healthy level of independent analysis, leading to groupthink. I do not have a good way to assess the level of groupthink in the third study, in part because avenues of independent thought are neither easy to categorize nor measure. One possible approach to avoid groupthink could be to design tools that promote discussion of alternative hypotheses [36]. Mechanisms to find and highlight disagreements in the findings might be good for initiating such discussions. Research into causes of groupthink and mechanisms to prevent it are an important area of future research. The much stronger performance of CLIP groups in our study suggests that the awareness benefits of LCW outweigh costs such as groupthink. However, further studies should look into this matter.

Chapter 7

Summary and Contributions

This thesis addresses the complexities of supporting the collaborative sensemaking process and challenges for designers of sensemaking support tools. Such tools should provide support for visualization as well as communication and coordination. Analysts need to organize their findings, hypotheses, and evidence, share their analytical findings with their collaborators, and coordinate their work activities. This thesis contributed a better understanding of collocated collaborative Visual Analytics processes and challenges. In addition I proposed design considerations to guide the development of collocated collaborative visual analytics tools.

In this chapter I briefly review the main contributions of the three phases of this research.

In the first phase of this research I conducted an observational study. The importance of record-keeping activities emerged from the results. The main contributions of the first phase include:

C1: A characterization of the collocated collaborative visual analytics process and activities during collocated collaborative VA, demonstrating the importance of record-keeping activities during this process.

C2: A categorization of notes based on their contents, scope, and usage. This rich understanding of note categories and usage could be beneficial for developers of collaborative VA tools.

C3: A set of guidelines to improve the design of record-keeping for collo-

cated collaborative visual analytics tools.
These design suggestions included: integration level for notes and visualization states, privacy considerations for individual notes as well as the ability to share them if desired, supporting record-keeping for different collaboration styles, note manipulation and management options (such as highlighting, grouping, and summarizing notes), and discussions around input mechanisms to support note taking on shared displays.

In the second phase, I conducted another user study which demonstrated how record-keeping support can be integrated into a tool to support collocated collaborative VA. The main contributions of this phase were:

C4: Validated the importance of externalization support during collaborative VA and revealed new ways to provide externalization support.

C5: A characterization of users' actions on visual record-keeping as well as their key intentions for each action.
Derived actions, intentions and also their dependence on the analytical phase and collaboration style. These could be used as a fundamental base to improve the design of record-keeping tools.

C6: Design guidelines for collaborative VA tools to improve record-keeping support.
These guidelines included providing various views of recorded material, showing manually saved rather than automatically saved items by default, enabling people to review collaborators' work unobtrusively, and automatically recommending items related to a user's analytical task.

In the third phase, I proposed, implemented and evaluated a new method, LCW (Linked Common Work), to increase awareness among team members. The LCW technique significantly improved group performance and communication and coordination within the groups. This method closely resembled *collaborative brushing and linking* [45] in which certain actions of one collaborator on a visualization are visible to other collaborators through their own views. However, I extended this concept to support the sensemaking loop. Then through an in depth quantitative analysis along with qualitative evaluation, I evaluated the effectiveness of this technique. This study demonstrated that LCW can increase awareness among team members. Furthermore, after summarizing and discussing the results from this user study, I proposed a theoretical collaboration model derived from the user study results. The key contributions were:

C7: Introduction, implementation and evaluation of LCW techniques in the context of collaborative sensemaking.

C8: Design guidelines to embed collaborative thinking spaces into collaborative VA tools.
In particular, I suggested how combining CLIP with a complementary tool that automatically identifies entities or relationships between documents could further increase awareness and support team collaboration. Furthermore, I advised adding visual indications that distinguish past work from current changes, adding a list of new changes, and using visual cues such as colour saturation to indicate a node's age. Also suggested adding a summary evidence file to help with publication of the results. Finally, I proposed some suggestions for scalability.

C9: A set of metrics to measure awareness, coordination and communication for collaborative VA.

Overall this thesis revealed that sharing externalizations (i.e. any recorded information related to the course of analysis such as insights, hypotheses, to-do lists or reminders for further analysis) could increase awareness and assist team members to better communicate and coordinate their work activities. I contributed an understanding of how analysts use VA tools during collocated collaboration. Through a series of observational user studies, I investigated how we can better support this complex process. More specifically, I empirically studied recording and sharing of analytical results. To this end, I implemented and evaluated two prototype tools to be able to understand the effects of these tools on collaboration mechanics. These user studies along with various literature surveys on each specific topic resulted in a collection of guidelines to support recording and sharing externalizations. In addition, I proposed, implemented and evaluated several mechanisms to increase awareness among team members, resulting in more effective coordination and communication during the collaborative sensemaking process.

However, the main contributions of this thesis are (1) identification and characterization of record-keeping behaviours as an important component of visual analytics and (2) introduction and evaluation of LCW technique to provide evidence in support of automatically identifying and presenting linked common work.

Appendices

Appendix A

Materials for the Observational Study

A.1 Consent Form

University of Victoria, Department of Computer Science

Participant Consent Form

Participant ID:

Human Factors in Visualization

You are being invited to participate in a study entitled [Collaborative visualization around large displays] that is being conducted by Melanie Tory and Narges Mahyar. Melanie Tory is a faculty member in the department of Computer Science at the University of Victoria and you may contact her if you have further questions by email at mtory@cs.uvic.ca or by phone at (250) 472-5798. This research is funded by NSERC.

The purpose of this research project is to investigate how people use large interactive surface to collaboratively explore and analyze data. Research of this type is important because it allows us to design better data tools to allow more effective, efficient, and enjoyable analysis of data in a variety of applications.

If you agree to voluntarily participate in this research, your participation will include:
o Filling out a background questionnaire that asks about your experience with computer technology and data analysis applications, as well as personal characteristics such as age and gender.
o Completing computer-based tasks.
o Participating in a verbal interview.
o Filling out a questionnaire about the computer-based tasks and tools you experienced.
o Being video and audio-taped.
o Being watched by live observers.
The research session take place at [Uvic ECS 654].

There are no known or anticipated risks to you by participating in this research.

Your participation in this research must be completely voluntary. If you do decide to participate, you may withdraw at any time without any consequences or any explanation. If you do withdraw, we will ask whether we may use your data for data analysis. If you decline, your data will be destroyed.

Your confidentiality and the confidentiality of the data will be protected by identifying data only with a participant number rather than your name, password-protecting computer files, and storing video and audio tapes in a locked office. Because this is a group study, confidentiality cannot be fully guaranteed since other participants in your group may know your identity. Confidentiality may not be guaranteed because the nature or size of the sample from which participants are drawn may make it possible to identify individual participants.

It is anticipated that the results of this study will be shared with others in the following ways:
o Published articles
o Conference presentations
o Video publications
o Theses
o Internet project descriptions

Data from this study will be disposed of within 5 years. Electronic data will be erased, paper copies will be shredded, and video/audio tapes will be recorded over or physically destroyed.

In addition to being able to contact the researcher at the above phone numbers, you may verify the ethical approval of this study, or raise any concerns you might have, by contacting the Associate Vice-President, Research at the University of Victoria (250-472-4545).

Your signature below indicates that you understand the above conditions of participation in this study and that you have had the opportunity to have your questions answered by the researchers.

Name of Participant ________________Signature________________Date________________

Email address:____________ Gender: F M Age: ____________ Degree:____________

A copy of this consent will be left with you, and a copy will be taken by the researcher.

A.2 Introduction

Introduction to Explorer

First you can upload data from a spreadsheet and choose the automatic mode. Different type of charts:

1. Comparison
2. Contribution
3. Correlation: You can activate it by using second measure
4. Trend

Explorer considers: Measures as X axis and dimension as Y axis. For example, you would like to understand who are the best players of all time. You choose point as measure and name as dimension. You can see your chart in different ways. Below is a step by step tutorial to find out who are the best players of all time.

1. Look at both points and games played, choosing game score as a second measure. You will see that Wayne and Mario are the best.
2. Show Correlation chart (now you have activated correlation chart)
3. Who has the highest score?
4. Players born in 1974? (The search option helps you to answer this question, just type the year in search box)
5. Hometown hero? For example, search Burnaby, select Burnaby as a 2nd measure and name as dimension.

Also I will show you how to make all types of charts and how to save or send a chart.

A.3 Task

Suppose you are a team of managers in a retail chain, and you want to plan out future marketing campaigns. You have a spreadsheet of sales data for the past few years and you teamed up to get a reasonable overview of the company's sales across regions, products, and time.

Part I: Warm-up Questions
As a team, find answers to the following questions.

1. What is the trend in sales revenue over 2003?

2. How does the 2003 margin compare to previous years?

3. What are the trends in the timing of purchases of Jackets?

4. Compare quantity of sales in New York, Chicago, and San Francisco.

5. Does sales revenue correlate with quantity sold for all products? If not, which products are inconsistent?

6. Q3 has lower sales revenue than the other quarters. Is this Q3 low consistent across all products? All cities?

Part II: Scenario

Please note that a different copy of the task was printed for each of the three participants in each group. They were identical with one single difference: the name of the state to be responsible for. Below is a sample scenario for the state of Texas:

You are the manager for Texas. The other two participants are managers for New York and California.

Together, you have a maximum marketing budget of $100,000. You want to make a recommendation to the CEO about how to allocate this budget among the three states for 2004. Based on the information available in the dataset, negotiate and allocate the budget among the three states. Include a breakdown of the budget by product if that seems reasonable. Prepare a presentation for the CEO to justify the reasoning behind your budget allocation.

When allocating the marketing budget, you may wish to consider information such as:
- History of sales in each region, for each product
- Popularity of different products in different regions
- Population of consumers in each state

States	**Approximate Population (millions)**
California	36.5
New York	19.3
Texas	23.9

Table A.1: Approximate population of Three States

Remember that you represent one state. You should try to obtain as much of the budget as you reasonably can for your state, while still cooperating to ensure the success of the company as a whole.

A.4 Follow up Interview

-How familiar are you in your understanding of the data set?

-What do you think is the reason why you got stuck at this point?

-What information would have helped you to overcome this problem?

-What features might have helped you to overcome this problem?

-What other comments do you have on this study?

Appendix B

Materials for the CoSpaces Study

B.1 Consent Form

University of Victoria, Department of Computer Science

Participant Consent Form

Participant ID:

Human Factors in Visualization

You are being invited to participate in a study entitled [Collaborative visualization around large displays] that is being conducted by Melanie Tory, Narges Mahyar, and Ali Sarvghad. Melanie Tory is a faculty member in the department of Computer Science at the University of Victoria and you may contact her if you have further questions by email at mtory@cs.uvic.ca or by phone at (250) 472-5798. This research is being funded by NSERC.

The purpose of this research project is to investigate how people use information, and how different visual representations of data affect how people perform tasks such as data analysis and decision making. Research of this type is important because it allows us to design better data displays to allow more effective, efficient, and enjoyable analysis of data in a variety of applications..

If you agree to voluntarily participate in this research, your participation will include:
o Filling out a background questionnaire that asks about your experience with computer technology and data analysis applications, as well as personal characteristics such as age and gender.
o Completing computer-based tasks.
o Participating in a verbal interview.
o Filling out a questionnaire about the computer-based tasks and tools you experienced.
o Being video and audio-taped.
o Being watched by live observers.
The research session take place at [Uvic ECS 654].

There are no known or anticipated risks to you by participating in this research.

Your participation in this research must be completely voluntary. If you do decide to participate, you may withdraw at any time without any consequences or any explanation. If you do withdraw, we will ask whether we may use your data for data analysis. If you decline, your data will be destroyed.

Your confidentiality and the confidentiality of the data will be protected by identifying data only with a participant number rather than your name, password-protecting computer files, and storing video and audio tapes in a locked office. Because this is a group study, confidentiality cannot be fully guaranteed since other participants in your group may know your identity. Confidentiality may not be guaranteed because the nature or size of the sample from which participants are drawn may make it possible to identify individual participants.

It is anticipated that the results of this study will be shared with others in the following ways:
o Published articles
o Conference presentations
o Video publications
o Theses
o Internet project descriptions

Data from this study will be disposed of within 5 years. Electronic data will be erased, paper copies will be shredded, and video/audio tapes will be recorded over or physically destroyed.

In addition to being able to contact the researcher at the above phone numbers, you may verify the ethical approval of this study, or raise any concerns you might have, by contacting the Associate Vice-President, Research at the University of Victoria (250-472-4545).

Your signature below indicates that you understand the above conditions of participation in this study and that you have had the opportunity to have your questions answered by the researchers.

Name of Participant ________________ Signature ________________ Date ________________

Email address: ______________ Gender: F M Age: ______________ Degree: ______________

A copy of this consent will be left with you, and a copy will be taken by the researcher.

B.2 Introduction

We have implemented an application for co-located collaborative visual analysis. The application that you will see is a prototype that is in the early stages of development and now we want to get feedback from users before the design is too far along.

The objective of this study is to evaluate our design and implementation decisions such as our proposed workspace, note taking and visualization history mechanisms. We also aim to evaluate usability of the system for co-located work, and to understand how you use any saved information.

I want to stress that we really want your honest feedback about the strengths and weaknesses of this program. Knowing what works and what doesn't work will help us to improve the next version.

Functionalities to be explained to the participants:
1- How to create a new workspace (also multiple workspaces)
2- How to create a chart (selecting chart type and mapping and filtering)
3- Talk about the data
4- How to map variables, how to undo mapping, target highlighting
5- How to filter data, how to undo filtering
6- How to clear a chart
7- How to take a note, text wrapping, explain colour coding of the notes and scrolling
8- How to review a note
9- How to reuse history, scrolling of history, reloading
10- How to use tabs, view other's current chart, reload from their history

Interactions to be explained to the participants:
11- Rotate
12- Zoom
13- Pan
14- Drag and drop

Explain where to touch to perform zoom and pan on keyboard and workspace

If you have questions or require assistance, please ask the observer. However, we would like you to first attempt to complete each task with the help of your partner.

Note that this study includes two tasks. In the first task you have some focused questions which will help you become familiar with the system as well as the dataset. The second task is an open-ended scenario that requires making a decision and creating a final report.

B.3 Task

Group:
Participant ID:
Date:
Task 1
Partcipant 1: NY Please follow all steps that are listed. We encourage you to make comments and discuss things with your partner aloud, as this will help us uncover problems with the prototype.

Suppose you are a team of managers in a retail chain, and you want to plan out future marketing campaigns. You have a spreadsheet of sales data for three years and you teamed up to get a reasonable overview of the company's sales across regions, products, and time. Please complete the following tasks:

1. Find out the trend in sales revenue over 2003.
a. Create a workspace.
b. Select line chart from the chart type menu.
c. Drag and drop sales revenue and quarter to yellow and blue highlighted areas and filter it based on 2003.
d. The current chart will appear in the history bar when you create a new chart.

2. Create a bar chart to see quantity of sales over category in NY.

3. How does the 2003 margin compare to previous years (2001 & 2002)?
a. Create a bar chart showing margin over year and compare their values.

4. What are the trends in the timing of purchases of Jackets and shirt-waist?
a. Create a bar chart that shows Quantity sold over year. Filter Lines to see data for Jackets.
b. Record your findings as a note.
c. Remove Jackets and filter based on shirt-waist.
d. Take a note.

5. Compare quantity of sales of different categories in NY and CA.

a. You created a chart for NY (Your team member created the same chart for California).
b. Create a new workspace, click on your partner's tab to view his/her workspace. From the history bar, reload the chart created for Quantity sold of products in CA.
c. Bring the new workspace next to yours and compare your NY chart to your team member's CA chart.
d. Take a note of your findings.

Task 2

Assume you are a new financial data analyst of a company that sells clothing to customers in the US over the Internet. Following is a list of the most popular product lines:

o Sweat t shirt
o Shirt waist
o Accessories
o Dresses
o Sweaters
o Outwear

You should divide this task with your partner. You will look at the first three items (underlined) and your team member will look at the rest. Analyze the sales data (Sales Revenue, Margin, Quantity sold) of the last 3 years over time, state and cities for these three items.

At the end, you and your team member should discuss and share your findings to be prepared to report back to your CEO.

B.4 Questionnaire

1- Do you feel that you successfully completed all the tasks on the task sheet?
Yes No

2- In relation to other software I have used, I found this prototype to be:
Very difficult to use 1 2 3 4 5 6 Very easy to use

3- The menu items were well organized and functions were easy to find.
Strongly disagree 1 2 3 4 5 6 Strongly agree

4- I immediately understood the function of each menu item.
Strongly disagree 1 , 2 3 4 5 6 Strongly agree

5- I found navigating around the prototype screen to be:
Very difficult 1 2 3 4 5 6 Very easy

6- It was easy and intuitive to reuse saved charts.
Strongly disagree 1 2 3 4 5 6 Strongly agree

7- The note taking mechanism was easy to use.
Strongly disagree 1 2 3 4 5 6 Strongly agree

8-I found my recorded information (charts and notes) useful.
Strongly disagree 1 2 3 4 5 6 Strongly agree

9-I found my team member's recorded information (charts and notes) useful.
Strongly disagree 1 2 3 4 5 6 Strongly agree

10- My overall impression of the prototype is:
Very negative 1 2 3 4 5 6 Very positive

11- In your opinion, what are the most useful features of the prototype?

Comments (please write down the application's problems. Your constructive suggestions are also appreciated):

B.5 Follow up Interview

1. At what point did you get stuck? What was the reason? What are the most confusing features of the prototype?

2. What are the most useful features?

3. How did you find history items? Did they help you in any way?

4. How about note taking? Did you find it useful?

5. What are your suggestions to improve the application?

6. What other comments do you have on this study?

Appendix C

Materials for the CLIP Study

C.1 Consent Form

University of Victoria, Department of Computer Science

Participant Consent Form

Participant ID:

Human Factors in Visualization

You are being invited to participate in a study entitled [Collaborative visualization around large displays] that is being conducted by Melanie Tory and Narges Mahyar. Melanie Tory is a faculty member in the department of Computer Science at the University of Victoria and you may contact her if you have further questions by email at mtory@cs.uvic.ca or by phone at (250) 472-5798. This research is being funded by NSERC.

The purpose of this research project is to investigate the effect of providing a sharing mechanism during collaborative analysis. Research of this type is important because it allows us to design better data displays to allow more effective, efficient, and enjoyable analysis of data in a variety of applications.

If you agree to voluntarily participate in this research, your participation will include:
o Filling out a background questionnaire that asks about your experience with computer technology and data analysis applications, as well as personal characteristics such as age and gender.
o Completing computer-based tasks.
o Participating in a verbal interview.
o Filling out a questionnaire about the computer-based tasks and tools you experienced.
o Being video and audio-taped.
o Being watched by live observers.
The research session take place at [Uvic ECS 654].

There are no known or anticipated risks to you by participating in this research.

Your participation in this research must be completely voluntary. If you do decide to participate, you may withdraw at any time without any consequences or any explanation. If you do withdraw, we will ask whether we may use your data for data analysis. If you decline, your data will be destroyed.

Your confidentiality and the confidentiality of the data will be protected by identifying data only with a participant number rather than your name, password-protecting computer files, and storing video and audio tapes in a locked office. Because this is a group study, confidentiality cannot be fully guaranteed since other participants in your group may know your identity. Confidentiality may not be guaranteed because the nature or size of the sample from which participants are drawn may make it possible to identify individual participants.

It is anticipated that the results of this study will be shared with others in the following ways:
o Published articles
o Conference presentations
o Video publications
o Theses
o Internet project descriptions

Data from this study will be disposed of within 5 years. Electronic data will be erased, paper copies will be shredded, and video/audio tapes will be recorded over or physically destroyed.

In addition to being able to contact the researcher at the above phone numbers, you may verify the ethical approval of this study, or raise any concerns you might have, by contacting the Associate Vice-President, Research at the University of Victoria (250-472-4545).

Your signature below indicates that you understand the above conditions of participation in this study and that you have had the opportunity to have your questions answered by the researchers.

Name of Participant ________________ Signature________________ Date________________

Email address:______________Gender: F M Age: ______________Degree:______________

A copy of this consent will be left with you, and a copy will be taken by the researcher.

C.2 Introduction

C.2.1 Introduction for Baseline Groups

The dataset is a mixed of real and fictitious information. It contains a scenario, newsletters and some mixed materials. We will start with a tutorial, then will give you 90 minutes to solve the task, and will end the study with an interview. We will give you a 5 minute break in the middle of the session. This mystery cannot be solved individually in 90 minutes, but if you work together there is a great chance to be able to solve it. There is a prize for the team with the most correct results. We would like to see how our tool helps you to collaboratively work together. So please use the tool even your results are not complete. You are allowed to talk to each other during the study.

System's Features and Interactions

1-Create a project

2-Create a node
o Add Name
o Add Time (optional)
o Add a note (Optional: set note scope (public/private), default: public)
o Attach evidence (Mandatory: adding at least one evidence)

Interactions with a **node**:
o Show details
o Edit
o Delete
o Link to a node
o Flag
o Filter: people, location, organization, flagged

Interactions with a **note**:
o Edit
o Delete
o Close

o Open externally
o Sort: time, type, owner

3- Create a link
o Add description
o Add note (optional)
o Attach evidence (optional)

Interactions with a **Link**:
o Show details
o Edit
o Delete

4- Collaborative features:
o **Tabs**: view others' work through tabs
o Collaborator's node: add, show detail
o Collaborator's note and link: anything but edit and delete

5- Evidence Cloud interaction:
o Click on a document to see where and who has attached it as evidence.

C.2.2 Introduction for CLIP Groups

The dataset is a mixed of real and fictitious information. It contains a scenario, newsletters and some mixed materials. We will start with a tutorial, then will give you 90 minutes to solve the task, and will end the study with an interview. We will give you a 5 minute break in the middle of the session. This mystery cannot be solved individually in 90 minutes, but if you work together there is a great chance to be able to solve it. There is a prize for the team with the most correct results. We would like to see how our tool helps you to collaboratively work together. So please use the tool even your results are not complete. You are allowed to talk to each other during the study.

System's Features and Interactions

1-Create a project

2-Create a node
o Add Name
o Add Time (optional)
o Add a note (Optional: set note scope (public/private), default: public)
o Attach evidence (Mandatory: adding at least one evidence)

Interactions with a **node**:
o Show details
o Edit
o Delete
o Link to a node
o Flag
o Filter: people, location, organization, flagged

Interactions with a **note**:
o Edit
o Delete
o Close
o Open externally

o Sort: time, type, owner

3- Create a link
o Add description
o Add note (optional)
o Attach evidence (optional)

Interactions with a **Link**:
o Show details
o Edit
o Delete

4- Collaborative features:
o **Combine all projects in your view**
o **Combine projects with common work**
o **Tabs**: view others' work through tabs
o Collaborator's node: add, show detail
o Collaborator's note and link: anything but edit and delete

5- Evidence Cloud interaction:
o Click on a document to see where and who has attached it as evidence.

C.3 Scenario

Stegosaurus Threat Stream Data Set Scenario

Another Tale of Alderwood, Washington

Note: This data set contains fictitious information. No part of this dataset should be taken as real or considered in any real analytical process. This data is for testing and evaluation of analytic tools.

Welcome to Alderwood, Washington, a fictitious town in central Washington state.

Already reeling from revelations of corruption and conspiracy, the people of Alderwood are left wondering What next? The head of the Alderwood Police Department believes he has a hint toward the answer to that question.

An incident from early 2004 (see The Alderwood News article from 1/16/2004 entitled Intoxicated Driver Makes Unexpected Deposit at Alderwood National Bank) has been nagging at him since it happened. Now, with the FBI's attention already focused on Alderwood, he has decided the time has come to shed some light on the situation.

This incident still doesn't sit right with me. Despite initial reports, no alcohol or drugs were found in the man's blood. That stuff he was spouting sounded like something you'd hear from a cult or something. If not for all the recent goings on, I might think I was being paranoid. I want this looked into.

The objective: To identify evidence, generate hypotheses and formulate next-step action suggestions for additional information-gathering. Your results should include a diagram and explanation of the social network, hypotheses (there could be more than one) of what is going on, and suggestions for the field officers to gather additional supporting evidence. The evidence presented should support your assertions of what the roles and relationships are and the motivations of the person(s) involved. Good luck!

The data includes:

1-Scenario: the scenario and the status of the investigation

2-Alderwood Daily Newsletters: news stories from the Alderwood Daily News plus a few other items collected by the previous investigators

3-Map: A map of Alderwood and vicinity

4-Mixed: Including:

-A spreadsheet about one of Alderwood's leading industries (apple orchards)

-Images: Ads taken from the newspaper

-Missile Silo document

-Dangerous Chemicals document

-South American Terrorist Organization document

C.4 Follow Up Interview

C.4.1 Follow Up Group Interview for Baseline Groups

Please note that questions for both conditions were very similar, the only major difference was that CLIP groups had two extra questions (shown in bold font).

1. What are the most useful features of CLIP? (List 3)
2. What are the three most confusing features?
3. At what point did you start to collaborate closely?
4. Which features encouraged you to collaborate closely?
5. How did you find adding evidence to each node?
6. How did you find adding notes to each node?
7. How did you find tabs?
8. Were you aware of what your teammates were doing most of the time?
9. Was the tool helpful in avoiding redundant work?
10. What are your suggestions to improve the system? Or comments on the study?
11. Comments on layout?

C.4.2 Follow Up Group Interview for CLIP Groups

1. What are the most useful features of CLIP? (List 3)
2. What are the three most confusing features?
3. At what point did you start to collaborate closely?
4. Which features encouraged you to collaborate closely?
5. **How did you find combining common items? (useful, confusing, distracting)**
6. **Did you combine all work or only work with common entities? How many times did you find critical information through others' work?**
7. How did you find adding evidence to each node?
8. How did you find adding notes to each node?
9. How did you find tabs?
10. Were you aware of what your teammates were doing most of the time?
11. Was CLIP helpful in avoiding redundant work?
12. What are your suggestions to improve the system? Or comments on the study?
13. Comments on layout?

Bibliography

[1] Christopher Andrews, Alex Endert, and Chris North. Space to think: large high-resolution displays for sensemaking. In *Proc. SIGCHI Conference on Human Factors in Computing Systems*, pages 55–64. ACM, 2010.

[2] Aruna D. Balakrishnan. *Sensemaking with Shared Visualizations: Investigating the Effects of Visualizations in Remote Collaborative Analysis*. PhD Thesis, Carnegie Mellon University, USA, 2011.

[3] Aruna D. Balakrishnan, Susan R. Fussell, and Sara Kiesler. Do visualizations improve synchronous remote collaboration? In *Proc. SIGCHI Conference on Human Factors in Computing Systems*, pages 1227–1236. ACM, 2008.

[4] Aruna D. Balakrishnan, Susan R. Fussell, Sara Kiesler, and Aniket Kittur. Pitfalls of information access with visualizations in remote collaborative analysis. In *Proc. ACM conference on Computer supported cooperative work*, pages 411–420. ACM, 2010.

[5] Louis Bavoil, Steven P. Callahan, Patricia J. Crossno, Juliana Freire, Carlos E. Scheidegger, Cláudio T. Silva, and Huy T. Vo. VisTrails: Enabling interactive multiple-view visualizations. In *Visualization, 2005. VIS 05. IEEE*, pages 135–142. IEEE, 2005.

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

[7] Francoise Boch and Annie Piolat. Note taking and learning: A summary of research. *The WAC Journal*, 16:101–113, 2005.

[8] Lauren Bradel, Alex Endert, Kristen Koch, Christopher Andrews, and Chris North. Large high resolution displays for co-located collaborative sensemaking: Display usage and territoriality. *International Journal of Human-Computer Studies*, 2013.

[9] Susan E. Brennan, Klaus Mueller, Greg Zelinsky, I. Ramakrishnan, David Scott Warren, and Arie Kaufman. Toward a multi-analyst, collaborative framework for visual analytics. In *Proc. Visual Analytics Science And Technology (VAST)*, pages 129–136. IEEE, 2006.

[10] Burke H. Bretzing and Raymond W. Kulhavy. Notetaking and depth of processing. *Contemporary Educational Psychology*, 4(2):145–153, 1979.

[11] Michael Broughton, Jeni Paay, Jesper Kjeldskov, Kenton O'Hara, Jane Li, Matthew Phillips, and Markus Rittenbruch. Being here: designing for distributed hands-on collaboration in blended interaction spaces. In *Proc. 21st Annual Conference of the Australian Computer-Human Interaction Special Interest Group: Design: Open 24/7*, pages 73–80. ACM, 2009.

[12] Stuart K Card, Jock D Mackinlay, and Ben Shneiderman. *Readings in information visualization: using vision to think*. Morgan Kaufmann, 1999.

[13] Yang Chen, Jamal Alsakran, Scott Barlowe, Jing Yang, and Ye Zhao. Supporting effective common ground construction in asynchronous collaborative visual analytics. In *Visual Analytics Science and Technology (VAST), 2011 IEEE Conference on*, pages 101–110. IEEE, 2011.

[14] Yang Chen, Scott Barlowe, and Jing Yang. Click2annotate: Automated insight externalization with rich semantics. In *Visual Analytics Science and Technology (VAST), 2010 IEEE Symposium on*, pages 155–162. IEEE, 2010.

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

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

[17] Herbert H Clark and Susan E Brennan. Grounding in communication. *Perspectives on socially shared cognition*, 13(1991):127–149, 1991.

[18] Kristin A. Cook and James J. Thomas. Illuminating the path: The research and development agenda for visual analytics. Technical report, Pacific Northwest National Laboratory (PNNL), Richland, WA (US), 2005.

[19] John W Creswell and Vicki L Plano Clark. *Designing and conducting mixed methods research*. Wiley Online Library, 2007.

[20] Steve Easterbrook, Janice Singer, Margaret-Anne Storey, and Daniela Damian. Selecting empirical methods for software engineering research. In *Guide to advanced empirical software engineering*, pages 285–311. Springer, 2008.

[21] Micheline Elias and Anastasia Bezerianos. Annotating bi visualization dashboards: Needs & challenges. In *Proc. 2012 ACM annual conference on Human Factors in Computing Systems*, pages 1641–1650. ACM, 2012.

[22] Sean E. Ellis and Dennis P. Groth. A collaborative annotation system for data visualization. In *Proc. working conference on Advanced visual interfaces*, pages 411–414. ACM, 2004.

[23] C. Fassnacht and D. Woods. Transana (version 2.20)[computer software]. *Madison, WI: School of Education at the University of Wisconsin-Madison*, 2007.

[24] Clifton Forlines, Alan Esenther, Chia Shen, Daniel Wigdor, and Kathy Ryall. Multi-user, multi-display interaction with a single-user, single-display geospatial application. In *Proc. 19th annual ACM symposium on User interface software and technology*, pages 273–276. ACM, 2006.

[25] Clifton Forlines and Chia Shen. Dtlens: multi-user tabletop spatial data exploration. In *Proc. 18th annual ACM symposium on User interface software and technology*, pages 119–122. ACM, 2005.

[26] Vagner Luiz Gava, Mauro de Mesquita Spinola, Antonio Carlos Tonini, and José Cardenas Medina. The 3c cooperation model applied to the classical requirement analysis. *JISTEM-Journal of Information Systems and Technology Management*, 9(2):235–264, 2012.

[27] Victor González and Alfred Kobsa. Benefits of information visualization systems for administrative data analysts. In *Information Visualization, 2003. IV 2003. Proceedings. Seventh International Conference on*, pages 331–336. IEEE, 2003.

[28] Carsten Görg, Zhicheng Liu, and John Stasko. Reflections on the evolution of the visual analytics system. *Information Visualization*, 2013.

[29] David Gotz and Michelle X. Zhou. Characterizing users' visual analytic activity for insight provenance. *Information Visualization*, 8(1):42–55, 2009.

[30] David Gotz, Michelle X Zhou, and Vikram Aggarwal. Interactive visual synthesis of analytic knowledge. In *Visual Analytics Science And Technology, 2006 IEEE Symposium On*, pages 51–58. IEEE, 2006.

[31] Saul Greenberg. 1 observing collaboration: Group-centered design. 2008.

[32] Carl Gutwin and Saul Greenberg. Workspace awareness for groupware. In *Conference Companion on Human Factors in Computing Systems*, pages 208–209. ACM, 1996.

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

[34] Amir Hossein Hajizadeh, Melanie Tory, and Rock Leung. Supporting awareness through collaborative brushing and linking of tabular data. *Visualization and Computer Graphics, IEEE Transactions on*, 19(12):2189–2197, 2013.

[35] James Hartley. Notetaking in non-academic settings: a review. *Applied Cognitive Psychology*, 16(5):559–574, 2002.

[36] Jeffrey Heer and Maneesh Agrawala. Design considerations for collaborative visual analytics. *Information visualization*, 7(1):49–62, 2008.

[37] Jeffrey Heer, Jock Mackinlay, Chris Stolte, and Maneesh Agrawala. Graphical histories for visualization: Supporting analysis, communication, and evaluation. *Visualization and Computer Graphics, IEEE Transactions on*, 14(6):1189–1196, 2008.

[38] Jeffrey Heer, Frank van Ham, Sheelagh Carpendale, Chris Weaver, and Petra Isenberg. Creation and collaboration: Engaging new audiences for information visualization. In *Information Visualization*, pages 92–133. Springer, 2008.

[39] Jeffrey Heer, Fernanda B. Viégas, and Martin Wattenberg. Voyagers and voyeurs: Supporting asynchronous collaborative visualization. *Communications of the ACM*, 52(1):87–97, 2009.

[40] M. Shahriar Hossain, Christopher Andrews, Naren Ramakrishnan, and Chris North. Helping intelligence analysts make connections. In *Scalable Integration of Analytics and Visualization*, 2011.

[41] Seth Hunter, Pattie Maes, Stacey Scott, and Henry Kaufman. Memtable: an integrated system for capture and recall of shared histories in group workspaces. In *Proc. SIGCHI Conference on Human Factors in Computing Systems*, pages 3305–3314. ACM, 2011.

[42] Petra Isenberg. *Collaborative Information Visualization in Co-located Environments*. PhD Thesis, University of Calgary, CA, 2009.

[43] Petra Isenberg and Sheelagh Carpendale. Interactive tree comparison for co-located collaborative information visualization. *Visualization and Computer Graphics, IEEE Transactions on*, 13(6):1232–1239, 2007.

[44] Petra Isenberg, Niklas Elmqvist, Jean Scholtz, Daniel Cernea, Kwan-Liu Ma, and Hans Hagen. Collaborative visualization: definition, challenges, and research agenda. *Information Visualization*, 10(4):310–326, 2011.

[45] Petra Isenberg and Danyel 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.

[46] Petra Isenberg, Danyel Fisher, Meredith Ringel Morris, Kori Inkpen, and Mary Czerwinski. An exploratory study of co-located collaborative visual analytics around a tabletop display. In *Visual Analytics Science and Technology (VAST), 2010 IEEE Symposium on*, pages 179–186. IEEE, 2010.

[47] Petra Isenberg, Danyel Fisher, Sharoda A. Paul, Meredith Ringel Morris, Kori Inkpen, and Mary Czerwinski. Co-located collaborative visual analytics around a tabletop display. *Visualization and Computer Graphics, IEEE Transactions on*, 18(5):689–702, 2012.

[48] Petra Isenberg, Anthony Tang, and Sheelagh Carpendale. An exploratory study of visual information analysis. In *Proc. twenty-sixth annual SIGCHI conference on Human factors in computing systems*, pages 1217–1226. ACM, 2008.

[49] Tobias Isenberg, Petra Neumann, Sheelagh Carpendale, Simon Nix, and Saul Greenberg. Interactive annotations on large, high-resolution information displays. In *Proc. 2006 Conference Compendium of IEEE VIS, IEEE InfoVis, and IEEE VAST*, pages 124–125, 2006.

[50] TJ Jankun-Kelly, Kwan-Liu Ma, and Michael Gertz. A model and framework for visualization exploration. *Visualization and Computer Graphics, IEEE Transactions on*, 13(2):357–369, 2007.

[51] Robert Johansen. *Groupware: Computer support for business teams.* The Free Press, 1988.

[52] Nazanin Kadivar, Victor Chen, Dustin Dunsmuir, Eric Lee, Cheryl Qian, John Dill, Christopher Shaw, and R. Woodbury. Capturing and supporting the analysis process. In *Visual Analytics Science and Technology, 2009. VAST 2009. IEEE Symposium on*, pages 131–138. IEEE, 2009.

[53] Eser Kandogan, Juho Kim, Thomas P Moran, and Pablo Pedemonte. How a freeform spatial interface supports simple problem solving tasks. In *Proc. SIGCHI Conference on Human Factors in Computing Systems*, pages 925–934. ACM, 2011.

[54] Youn-ah Kang, C. Gorg, and John Stasko. Evaluating visual analytics systems for investigative analysis: Deriving design principles from a case study. In *Proc. Visual Analytics Science and Technology (VAST)*, pages 139–146. IEEE, 2009.

[55] Youn-ah Kang and John 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.

[56] Daniel Keim, Gennady Andrienko, Jean-Daniel Fekete, Carsten Görg, Jörn Kohlhammer, and Guy Melançon. *Visual analytics: Definition, process, and challenges.* Springer, 2008.

[57] Kenneth A. Kiewra, Stephen L. Benton, Sung-Il Kim, Nancy Risch, and Maribeth Christensen. Effects of note-taking format and study technique on recall and relational performance. *Contemporary Educational Psychology*, 20(2):172–187, 1995.

[58] KyungTae Kim, Waqas Javed, Cary Williams, Niklas Elmqvist, and Pourang Irani. Hugin: A framework for awareness and coordination in mixed-presence collaborative information visualization. In *ACM Intl. Conference on Interactive Tabletops and Surfaces*, pages 231–240. ACM, 2010.

[59] Robert E Kraut, Susan R Fussell, Susan E Brennan, and Jane Siegel. Understanding effects of proximity on collaboration: Implications for technologies to support remote collaborative work. *Distributed work*, pages 137–162, 2002.

[60] Olga Kulyk, Gerrit van der Veer, and Betsy van Dijk. Situational awareness support to enhance teamwork in collaborative environments. In *Proc. 15th European conference on Cognitive ergonomics: the ergonomics of cool interaction*, page 5. ACM, 2008.

[61] Heather Richter Lipford, Felesia Stukes, Wenwen Dou, Matthew E. Hawkins, and Remco Chang. Helping users recall their reasoning process. In *Visual Analytics Science and Technology (VAST), 2010 IEEE Symposium on*, pages 187–194. IEEE, 2010.

[62] Narges Mahyar, Ali Sarvghad, and Melanie Tory. A closer look at note taking in the co-located collaborative visual analytics process. In *Visual Analytics Science and Technology (VAST), 2010 IEEE Symposium on*, pages 171–178. IEEE, 2010.

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

[64] Narges Mahyar, Ali Sarvghad, Melanie Tory, and Tyler 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.

[65] Narges mahyar and Melanie Tory. Supporting communication and coordination in collaborative sensemaking. *Visualization and Computer Graphics, IEEE Transactions on*, 2014.

[66] Valérie Maquil, Thomas Psik, and Ina Wagner. The colortable: a design story. In *Proc. 2nd international conference on Tangible and embedded interaction*, pages 97–104. ACM, 2008.

[67] Gloria Mark, Keri Carpenter, and Alfred Kobsa. A model of synchronous collaborative information visualization. In *Information Visualization, 2003. IV 2003. Proceedings. Seventh International Conference on*, pages 373–381. IEEE, 2003.

[68] Gloria Mark and Alfred Kobsa. The effects of collaboration and system transparency on CIVE usage: an empirical study and model. *Presence: Teleoperators and Virtual Environments*, 14(1):60–80, 2005.

[69] Catherine C. Marshall. Work practices study: analysts and note taking. unpublished report. `http://www.csdl.tamu.edu/~marshall/oswrstudy.pdf`. Accessed: 2010-09-30.

[70] E Mcgrath. Methodology matters: Doing research in the behavioral and social sciences. In *Readings in Human-Computer Interaction: Toward the Year 2000 (2nd ed*. Citeseer, 1995.

[71] Will McGrath, Brian Bowman, David McCallum, Juan David Hincapié-Ramos, Niklas Elmqvist, and Pourang Irani. Branch-explore-merge: facilitating real-time revision control in collaborative visual exploration. In *Proc. ACM conference on Interactive tabletops and surfaces*, pages 235–244. ACM, 2012.

[72] Chii Meng, Motohiro Yasue, Atsumi Imamiya, and Xiaoyang Mao. Visualizing histories for selective undo and redo. In *Computer Human Interaction, 1998. Proceedings. 3rd Asia Pacific*, pages 459–464. IEEE, 1998.

[73] David R Millen, Audrey Schriefer, Diane Z Lehder, and Susan M Dray. Mind maps and causal models: using graphical representations of field research data. In *CHI'97 Extended Abstracts on Human Factors in Computing Systems*, pages 265–266. ACM, 1997.

[74] Thomas P Moran, William Van Melle, and Patrick Chiu. Spatial interpretation of domain objects integrated into a freeform electronic whiteboard. In *Proc. ACM Symp. User Interface Software and Technology*, pages 175–184. ACM, 1998.

[75] Meredith Ringel Morris, Jarrod Lombardo, and Daniel Wigdor. Wesearch: supporting collaborative search and sensemaking on a tabletop display. In *Proc. 2010 ACM conference on Computer supported cooperative work*, pages 401–410. ACM, 2010.

[76] Meredith Ringel Morris, Kathy Ryall, Chia Shen, Clifton Forlines, and Frederic Vernier. Beyond social protocols: Multi-user coordination policies for co-located groupware. In *Proc. 2004 ACM conference on Computer supported cooperative work*, pages 262–265. ACM, 2004.

[77] Dennis C. Neale, John M. Carroll, and Mary Beth Rosson. Evaluating computer-supported cooperative work: models and frameworks. In *Proc. 2004 ACM conference on Computer supported cooperative work*, pages 112–121. ACM, 2004.

[78] Kenton O'hara, Jesper Kjeldskov, and Jeni Paay. Blended interaction spaces for distributed team collaboration. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 18(1):3, 2011.

[79] Kenton O'Hara and Abigail Sellen. A comparison of reading paper and on-line documents. In *Proc. ACM SIGCHI Conference on Human factors in computing systems*, pages 335–342. ACM, 1997.

[80] Tom Oinn, Matthew Addis, Justin Ferris, Darren Marvin, Martin Senger, Mark Greenwood, Tim Carver, Kevin Glover, Matthew R. Pocock, Anil Wipat, et al. Taverna: a tool for the composition and enactment of bioinformatics workflows. *Bioinformatics*, 20(17):3045–3054, 2004.

[81] Kyoung S. Park, Abhinav Kapoor, and Jason Leigh. Lessons learned from employing multiple perspectives in a collaborative virtual environment for visualizing scientific data. In *Proc. third international conference on Collaborative virtual environments*, pages 73–82. ACM, 2000.

[82] Sharoda A. Paul and Madhu C. Reddy. Understanding together: sensemaking in collaborative information seeking. In *Proc. ACM Conf. on Computer Supported Cooperative Work*, pages 321–330. ACM, 2010.

[83] Peter Pirolli and Stuart Card. Information foraging. *Psychological review*, 106(4):643, 1999.

[84] Peter Pirolli and Stuart 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.

[85] Jakobsen Mikkel R. and Horanbæk Kasper. Collaborative work on a high-resolution multi-touch wall-display. *Trans. Computer-Human Interaction*, 21(2):340, 2014.

[86] Anthony C. Robinson. Collaborative synthesis of visual analytic results. In *Proc. Visual Analytics Science and Technology (VAST)*, pages 67–74. IEEE, 2008.

[87] Yvonne Rogers and Siân Lindley. Collaborating around vertical and horizontal large interactive displays: which way is best? *Interacting with Computers*, 16(6):1133–1152, 2004.

[88] Roy A. Ruddle, Rhys G. Thomas, Rebecca S. Randell, Philip Quirke, and Darren Treanor. Performance and interaction behaviour during visual search on large, high-resolution displays. *Information Visualization*, page 1473871613500978, 2013.

[89] Daniel M. Russell, Mark J. Stefik, Peter Pirolli, and Stuart K. Card. The cost structure of sensemaking. In *Proc. INTERACT'93 and CHI'93 conference on Human factors in computing systems*, pages 269–276. ACM, 1993.

[90] Sap business objects. explorer in the cloud. `https://create.ondemand.com/explorer`. Accessed: 2010-06-30.

[91] Ali Sarvghad, Narges Mahyar, and Melanie Tory. History tools for collaborative visualization. *Collaborative Visualization on Interactive Surfaces-CoVIS'09*, page 21, 2009.

[92] Mike Scaife and Yvonne Rogers. External cognition: how do graphical representations work? *International journal of human-computer studies*, 45(2):185–213, 1996.

[93] Stacey D. Scott, Karen D. Grant, and Regan L. Mandryk. System guidelines for co-located, collaborative work on a tabletop display. In *ECSCW 2003*, pages 159–178. Springer, 2003.

[94] Stacey D Scott, M. Sheelagh, T. Carpendale, and Kori M. Inkpen. Territoriality in collaborative tabletop workspaces. In *Proc. 2004 ACM conference on Computer supported cooperative work*, pages 294–303. ACM, 2004.

[95] Abigail J. Sellen and Richard Harper. *The myth of the paperless office*. The MIT Press, 2003.

[96] Ben Shneiderman. The eyes have it: A task by data type taxonomy for information visualizations. In *Visual Languages, 1996. Proceedings., IEEE Symposium on*, pages 336–343. IEEE, 1996.

[97] Yedendra Babu Shrinivasan and Jarke J. van Wijk. Supporting the analytical reasoning process in information visualization. In *Proc. SIGCHI conference on human factors in computing systems*, pages 1237–1246. ACM, 2008.

[98] Yedendra Babu Shrinivasan, D y, and Jie Lu. Connecting the dots in visual analysis. In *Proc. Visual Analytics Science and Technology (VAST)*, pages 123–130. IEEE, 2009.

[99] Georg Simmel. The number of members as determining the sociological form of the group. *The American Journal of Sociology*, 8(1):1–46, 1902.

[100] John Stasko, Carsten Görg, and Zhicheng Liu. Jigsaw: supporting investigative analysis through interactive visualization. *Information visualization*, 7(2):118–132, 2008.

[101] Tableau. `http://www.tableausoftware.com`.

[102] Anthony Tang, Melanie Tory, Barry Po, Petra Neumann, and Sheelagh Carpendale. Collaborative coupling over tabletop displays. In *Proc. SIGCHI conference on Human Factors in computing systems*, pages 1181–1190. ACM, 2006.

[103] Lucia Terrenghi, David Kirk, Abigail Sellen, and Shahram Izadi. Affordances for manipulation of physical versus digital media on interactive surfaces. In *Proc. SIGCHI conference on Human factors in computing systems*, pages 1157–1166. ACM, 2007.

[104] Matthew Tobiasz, Petra Isenberg, and Sheelagh Carpendale. Lark: Coordinating co-located collaboration with information visualization. *Visualization and Computer Graphics, IEEE Transactions on*, 15(6):1065–1072, 2009.

[105] Melanie Tory. User studies in visualization: A reflection on methods. In *Handbook of Human Centric Visualization*, pages 411–426. Springer, 2014.

[106] Melanie Tory, Sheryl Staub-French, Barry A Po, and Fuqu Wu. Physical and digital artifact-mediated coordination in building design. *Computer Supported Cooperative Work (CSCW)*, 17(4):311–351, 2008.

[107] Frédéric Vernier, Neal Lesh, and Chia Shen. Visualization techniques for circular tabletop interfaces. In *Proc. Working Conference on Advanced Visual Interfaces*, pages 257–265. ACM, 2002.

[108] Fernanda B. Viégas, Martin Wattenberg, Frank Van Ham, Jesse Kriss, and Matt McKeon. ManyEyes: a site for visualization at internet scale. *Visualization and Computer Graphics, IEEE Transactions on*, 13(6):1121–1128, 2007.

[109] Katherine Vogt, Lauren Bradel, Christopher Andrews, Chris North, Alex Endert, and Duke Hutchings. Co-located collaborative sensemaking on a large high-resolution display with multiple input devices. In *Human-Computer Interaction–INTERACT 2011*, pages 589–604. Springer, 2011.

[110] James R. Wallace, Stacey D. Scott, and Carolyn G. MacGregor. Collaborative sensemaking on a digital tabletop and personal tablets: prioritization, comparisons, and tableaux. In *Proc. SIGCHI Conference on Human Factors in Computing Systems*, pages 3345–3354. ACM, 2013.

[111] Colin Ware. *Information visualization: perception for design*. Elsevier, 2012.

[112] Karl E Weick. The collapse of sensemaking in organizations: The mann gulch disaster. *Administrative science quarterly*, pages 628–652, 1993.

[113] Karl E Weick. *Sensemaking in organizations*, volume 3. Sage, 1995.

[114] Mark A. Whiting, Chris North, Alex Endert, Jean Scholtz, Jereme Haack, Carrie Varley, and Jim Thomas. VAST contest dataset use in education. In *Proc. Visual Analytics Science and Technology (VAST)*, pages 115–122. IEEE, 2009.

[115] Wesley Willett, Jeffrey Heer, Joseph Hellerstein, and Maneesh Agrawala. Commentspace: structured support for collaborative visual analysis. In *Proc. SIGCHI Conference on Human Factors in Computing Systems*, pages 3131–3140. ACM, 2011.

[116] William Wright, David Schroh, Pascale Proulx, Alex Skaburskis, and Brian Cort. The sandbox for analysis: concepts and methods. In *Proc. SIGCHI conference on Human Factors in computing systems*, pages 801–810. ACM, 2006.

[117] Jiajie Zhang. The nature of external representations in problem solving. *Cognitive science*, 21(2):179–217, 1997.