


An Exploratory Study of Co-located Collaborative Visual Analytics Around a Tabletop Display

Petra Isenberg^{*,†} Danyel Fisher[†] Meredith Ringel Morris[†] Kori Inkpen[†] Mary Czerwinski[†]

^{*}INRIA

[†]Microsoft Research



(a) Analysts begin by searching for data. (b) A document is pulled out of a search result list. (c) A document is zoomed to be read and analyzed. (d) The workspace can be flexibly arranged and shared.

Figure 1: Several different aspects of the Cambiera system in use during a document analysis task.

ABSTRACT

Co-located collaboration can be extremely valuable during complex visual analytics tasks. This paper presents an exploratory study of a system designed to support collaborative visual analysis tasks on a digital tabletop display. Fifteen participant pairs employed Cambiera, a visual analytics system, to solve a problem involving 240 digital documents. Our analysis, supported by observations, system logs, questionnaires, and interview data, explores how pairs approached the problem around the table. We contribute a unique, rich understanding of how users worked together around the table and identify eight types of collaboration styles that can be used to identify how closely people work together while problem solving. We show how the closeness of teams' collaboration influenced how well they performed on the task overall. We further discuss the role of the tabletop for visual analytics tasks and derive novel design implications for future co-located collaborative tabletop problem solving systems.

Index Terms: K.5.3 [Information Interfaces and Presentation]: Group and Organization Interfaces

1 INTRODUCTION

Visual analytics tasks can have amorphous structure, ambiguous goals, and large amounts of data [14]. Many visual analytics problems, therefore, can be best solved by groups of analysts working together, face-to-face. For example, Chin et al. [5] found that intelligence analysts highly value the ability to collaborate face-to-face: *"It's when all of our analysts get together and work out the differences and challenge each other with facts that we get to a better and more prominent answer."* Despite these benefits, we have surprisingly few technologies to support co-located visual analytics: instead, teams typically share one screen or work on separate

computers, side by side. New technologies, such as interactive surfaces, offer increased screen real estate and multi-touch capabilities to build novel face-to-face collaborative workspaces. Yet, there have been few studies reporting on design and use of collaborative visual analytics workspaces for interactive surfaces. We try to fill this gap by providing a rich description of pairs working collaboratively, face-to-face, on a visual analytics task, using a tabletop system designed specifically for these tasks. We studied teams of two people solving the VAST 2006 Challenge, an intelligence analysis task that involves the exploration of a large text document collection. This task resembles many other scenarios in which pairs collaboratively attempt to solve complex problems over a large text corpus: for instance, lawyers making sense of a set of case files, military intelligence analysts trying to comb through millions of documents, co-authors researching old newspapers for help with a new story, or historians browsing through historic texts to bring past events into context. In these and many other tasks, analysts have considerable experience working face-to-face on traditional physical tables [5]. Our goal is to more fully understand the consequences of providing a digital tabletop analytics environment that incorporates the benefits of a shared space for synchronous collaborative work as well as computational support for searching and data sharing. Studying and formalizing the nature of collaborative analysis activities within such an environment is valuable to inform our understanding of digital tabletop visual analytics systems and provides valuable design advice for future systems, as we shall discuss. We begin by reviewing the Cambiera system [10], and then present a study and discuss our observations. Our goal is to address the digital table as a context for co-located visual analytics: to describe how pairs approached the problem on the table effectively; to discuss what features helped pairs with the analysis task, and to describe what additional features would be even more beneficial in supporting complex problem solving on the tabletop.

2 RELATED WORK

Our study is a rich examination of how teams communicate and coordinate face-to-face around a visual analytics tabletop system. While researchers have explored various tools for visual analytics support, our focus on co-located tabletop collaboration is unique

^{*}e-mail:petra.isenberg@inria.fr

[†]e-mail:{danyelf|merrie|kori|marycz}@microsoft.com

and revealed novel patterns of interacting that have not before been reported thoroughly. In this section, we discuss work that has inspired our own through a focus on collaborative work with complex analysis tasks and visualization support.

2.1 Distributed Collaborative Problem Solving

Most research projects exploring collaborative problem-solving have supported distributed, rather than face-to-face, collaboration. For example, Balakrishnan et al. [2] addressed the question of how teams can share visualizations remotely to solve a complex problem. They found that a visualization was most effective when both partners had full access to the shared visualization and could both synchronously interact with it; in a follow-up study [3], they further found that collaborators were more effective when they each had control over parts of the data. We include an implicitly shared context in our system, where we give each analyst independent access to all data items and data interactions on the table. Brennan et al. [4] take a slightly different approach by requiring explicit sharing and merging of data views during distributed analysis. This way, team members can work alone on a subtopic, and then switch to a shared view that supports joint sensemaking. Like our system, Paul and Morris' search-oriented CoSense [16] provides collaborators with awareness support: teammates can see each other's web search histories, as well as the documents that they each have found, in order to support collaborative information seeking. Keel's system [13] for distributed analysts provided awareness by analyzing information from a team member's workspace and suggesting relevant data to other remote analysts, allowing for implicit information sharing.

In contrast to these systems, ManyEyes [21] social visualizations connect people who may not have any previous relationship. People can asynchronously share data, create visualizations, and collaboratively explore, annotate, and comment on the visualizations. In contrast to these systems for distributed problem solving, our work addresses face-to-face collaboration around a shared display and fundamentally differs in terms of the workspace setup and the ability for team members to communicate quickly and effortlessly about the data on the surface. In addition, it allows team members to distinguish themselves in meaningful and useful ways. We included some of the features employed earlier for distributed systems—such as support for parallel work and joint search histories—in our system and study, as they seem to be valuable.

2.2 Co-located Collaborative Problem Solving

Tabletop technologies can allow teams to easily share tools and information face-to-face, and researchers are still learning how best to leverage tables for complex tasks. Research on collaborative visual analytics on tabletop displays is still in its infancy (e. g., [11]). We therefore discuss work from the Tabletop community that addresses more complex analysis tasks.

WeSearch [15] is related to our work in that it supports search, browsing, and sensemaking activities. Team members can share web searches, snippets, and annotations during their investigations. The DTLens system [6] demonstrates how focus+context techniques can be used for exploring large maps and diagrams on an interactive tabletop. DTLens addresses an important issue for collaborative problem solving: when multiple people simultaneously interact with information, parallel exploration may be hindered when interests between team members differ. DTLens resolves this with a notion of ownership, in which individuals can only edit their own views of the data. In contrast, CoTree, a collaborative tabletop system for the comparison and analysis of hierarchical data [9] allows collaborators to use multiple independent views and visualizations of the same dataset. These can be explored in parallel or brought together for comparison without any sense of ownership. WeSpace [22] addresses the issue of providing individual views, and parallel exploration by allowing researchers to bring their own visualiza-




Figure 2: Two search widgets. The top search widget (“bse”) is minimized. The bottom widget is maximized; each vertical bar represents a document that contains the word “fda.” The documents are ordered by publication date. The fourth and fifth documents also have green stripes and contain the word “bse.”




Figure 3: Overview of the system in use. Each person issues searches starting from their own search button in front of them (highlighted by circles). Touching the button brings up a keyboard (Figure 1a) to type search terms and issue the search. Here the left analyst issues blue, and the right issues orange searches.

tion applications on their own laptops to a multi-display environment. The table here is used to coordinate views from the different laptops on a wall around which shared viewing, discussion, and interpretation can happen. However, as the display is controlled by individual laptops, the system does not specifically facilitate simultaneous interaction with the shared views of the data. In our system we combined several of these approaches. We allowed collaborators to share data and work with multiple views of the same data source. We also made ownership explicit through color-coding but did not restrict view access. The next section describes these features in more detail. Finally, while a number of systems have been previously developed to support the complex tasks involved in intelligence work (e. g., [4, 13, 19]), to our knowledge there have been no studies of other systems in use in a tabletop context supporting visual analytics tasks.

3 CAMBIERA: TABLETOP SYSTEM DESCRIPTION

Our study examines pairs analyzing a shared document collection using Cambiera [10]. Cambiera is one of the first tabletop systems designed specifically for co-located collaborative visual analytics. It currently runs on the Microsoft Surface, a multi-touch tabletop display. As a tabletop system, Cambiera encourages analysts to face each other around the tabletop, while analyzing large text document collections collaboratively. Cambiera features four fundamental operations:

Search for documents: Each team member starts their analysis with Cambiera by searching for relevant documents by calling one of four on-screen keyboards from the search button in front of them (Figure 1a and Figure 3). Each keyboard specific to and oriented towards just one analyst returns a color-coded search result widget which holds virtual representations of the documents returned (Figure 2). Frequency bars under each document represent the total number of times the respective search term occurs in the document. The color hue of the widget is determined by the team member who issued the search. While hues between search words are only subtly different (and each analyst may not remember them in detail), each person is vividly distinct; in addition, the analyst can touch a document to see what words are associated with it.

Search result exploration: Each search widget can be expanded to see the returned list of documents ordered by their publication date (Figure 2). By running a finger across the document representations, detail-on-demand information is shown for each document. This information includes who has previously issued the search, what other search terms (regardless of searcher) found the same document, and which sentences contained the search term.

Document analysis: documents can be pulled out of search result lists (Figure 1b) and then freely placed in the workspace. By performing a zooming gesture the document's text can be accessed within a document reader (Figure 1c).

Workspace organization: All objects in Cambiera can be moved around the table, re-oriented, resized, or closed. Each analyst can place and stack documents and search results on the tabletop as they find most useful (Figure 1d). An analyst can also spin the documents around in order to push them to the other side, one of the several ways that Cambiera allows people to share documents.

One of the unique features of Cambiera is its support for co-located awareness. Each individual search widget is tagged with a unique, user-dependent color. Each document within the search widget in turn is annotated with colored stripes corresponding to the search terms that have been found within them. Because each search widget has its own color, the colored stripes on the search results indicate when they have been found by more than one word. In Figure 2, for example, one team member has searched for "FDA" (Federal Department of Agriculture); another has searched for "BSE" (Bovine Spongiform Encephalitis). The colored stripes indicate that the last two documents found by the search for "fda" also contain the word "bse." When documents are pulled out of a search result list, they maintain the colored bars that indicate which search generated them. In the document reader (Figure 1c), the text highlights all search terms which have found this document, regardless of which team member issued them. In addition, all states of a document include a specific icon (in the search result list, minimized in the workspace, and maximized in the reader) which indicates whether the document has been read, by whom and the shading of the document's background color gives a relative indication of read frequency. Since previous studies of collaborative data analysis have shown that team members frequently switch to phases of parallel work [12, 20], these awareness features were designed to encourage co-located team members to connect more closely with others and to be able to more easily synthesize their individual analysis results. We refer to the accompanying video for an overview of the Cambiera system in use.

4 USER STUDY

To design visual analytics tabletop tools for co-located collaborative work we need to understand how these systems are used by teams. In particular, we need a better understanding of how teams coordinate their activities over a tabletop and which specific system features are best suited to support effective co-located collaborative data analysis in this context. Our study was designed to provide a rich description of collaborative visual analysis on a tabletop dis-

play targeting this information need. In this exploration, we were guided by questions of understanding collaborative analytics on the tabletop.

- How do teams work together and coordinate their work with the tabletop?
- What patterns of collaboration do they carry out?
- How can their work patterns be used to inform the design of future interfaces for collaborative problem solving?

4.1 Participants

It is extremely difficult to obtain professional intelligence analysts for real-life intelligence analysis scenarios; therefore, we followed the procedures used by the VAST 2006 Challenge [7], which provides a baseline task that closely models an analysis scenario, but can be carried out by non-analysts. We describe the task further below. We recruited 15 pairs of people who were familiar with data analysis. Participants were required to have a Master's (or more advanced) degree, and to have self-reported as enjoying puzzles or mysteries. The members of each pair knew each other and had previously worked together in some form; subjects were co-workers, friends, family members, and married couples. Participants ranged in age from 25–55; ten couples were mixed-gender; three both women; two both male. On average participants spent 72 minutes ($stddev(\sigma) = 12$) working on the task.

4.2 The VAST 2006 Challenge

We based our experiment around the VAST 2006 Challenge, code-named "Stegosaurus" [7], a scenario that entails finding a hidden weapons-smuggling plot. Stegosaurus was developed by the National Visual Analytics Center at Pacific National Labs. We chose "Stegosaurus" because it is increasingly recognized as a standard visual analytics task [1]. The dataset is available to the public to download, and is usable by non-experts without training. Nonetheless, it is a complex task that requires participants to connect and infer facts buried in a set of several hundred documents. We used an extended version of the Stegosaurus document set containing 240 digital and 6 paper documents. Cambiera does not currently support search of non-textual materials like maps and images. Thus, we preloaded Cambiera with all of the newspaper and fact sheet articles but provided paper print-outs of the map and other images.

The task comes with a vague goal of figuring out a car crash; from there, the analyst needs to work through the dataset. No single document tells the whole story—of the 246 documents, just ten have true bearing on the story, while several others provide some background information but are not necessary, and the rest are irrelevant. Solving the task requires carefully filtering out irrelevant articles, reading articles in detail, and making several intuitive leaps. In between stories about fruit-picking season, one document tells us that people who ate at an apple festival got intestinal poisoning, probably from tainted meat. A second article says that a batch of apples stored in a silo was contaminated with nerve gas. It requires searching, reading, and intuition to recognize that the symptoms of nerve gas were likely mistaken for intestinal poisoning, and that they came from the same apples. Collaboration has the potential to improve the effectiveness of the analysis in several ways. Teams can cover more documents in the same amount of time, can discuss evolving hypotheses and contribute facts they may have found individually. At the same time, collaboration can be a mixed blessing if team members fail to share and/or synthesize their results. The goal of our evaluation is to study more closely how teams work together in solving this task with our tool, and to derive further requirements and guidelines for tabletop visual analytics tools.

4.3 Procedure

Our study used the VAST 2006 challenge data loaded in Cambiera. Participants were seated on opposite sides of the table, and re-

ceived a fifteen-minute tutorial on the features of Cambiera using a sample dataset, during which they were encouraged to experiment with the system and ask questions freely. They were read an introductory letter to the “Stegosaurus” problem, which explained the task context. That letter includes a starting clue that suggested a first document to read. Since our focus was on observing the pairs’ collaborative interactions with each other and with the system, and less on performance outcomes, we wanted to ensure that they were able to make progress in the task, and that they continued working. For this reason, we provided assistance to teams who did not progress in the task. The experimenter running the study was familiar with the dataset, and so was able to monitor teams’ progress. When teams stopped making progress, as judged by the experimenter (e.g., reading and re-reading distracter documents or reporting to be stuck) the experimenter provided assistance. During an assist, the experimenter did not provide new information, but rather asked the participants to clarify previous ideas that they had raised. We maintained consistency of assists by using only one experimenter, who followed a written protocol. Participants were also provided with pens and notepads for taking notes, and a small set of auxiliary images that are part of the “Stegosaurus” scenario. Participants reported their results verbally at the end of the study. We terminated the experiment when the team could produce a coherent story when asked for their hypotheses, and ended all experiments at one and a half hours. After the study, participants independently filled out a questionnaire (due to a technical error, three pairs did not fill out the questionnaire). The questionnaire elicited information on participant demographics and subjective awareness experience. Finally, the experimenter interviewed the pair to understand how they approached the problem and to get their feedback on the technology. Originally, we designed three experimental conditions that varied aspects of Cambiera’s visualization; in the different conditions, three different color schemes were used to subtly convey information about which documents one’s partner had read and which terms they had searched for. In a formal statistical analysis of the data we did not find significant differences between the number of searches performed, documents read, documents/searches passed between participants, or facts connected in each condition. Thus, we report observations that apply to all pairs, regardless of condition.

4.4 Data Analysis

In order to fully describe the event, we captured a variety of sources: each experimental session was video- and audio-recorded; in addition, the software captured a screenshot every minute and a time-stamped event log for interactions with the tabletop. One experimenter took notes in real time; a second experimenter engaged in a two-pass video coding in order to get a rich understanding of the ways in which teams solved the task. During the first video coding pass the following information was coded: collaboration styles based on a code set by Tang et al. [20], roles adopted by participants, use of awareness features, use of external information (shared or private notes), breakdowns or conflicts in group work, and how often workspace items were shared. After the first coding pass it became evident that the code set from Tang et al. [20] had to be extended to accommodate the different study situation. Whereas the original code set was developed for a situation in which participants shared the same representation, our participants could work with data, representations, and views of the data in parallel. Therefore, we extended the code set to more clearly distinguish when people shared views of the same data and when they were looking at the same data but using copies of the same information items and thus had different views of the same information. During the second coding pass, the video coder also took more detailed time-stamped notes on when participants switched to different types of collaboration styles using the extended code set as well as extended

notes on which facts and documents teams found as they progressed through the task. Written transcripts of participants’ answers to the post-session questionnaire were coded into higher-level categories including awareness (e.g., which information was missing or helpful), work styles (strategies, roles, sharing, and collaboration), and tool features commented on (liked and missing features). The log data was further parsed and quantitatively analyzed. Notes taken by participants and used during the experiment were also collected and analyzed. Together, the detailed analysis of video, field notes, system logs, and interview data provided a rich understanding of how participants solved the task, worked with Cambiera, and how they engaged with each other.

5 SOLVING A COMPLEX PROBLEM COLLABORATIVELY

With the exception of some research papers [1, 5] and the VAST contest reports, there are few descriptions of how users work on complex analysis tasks, and none of them discuss collaborative aspects in detail [18]. In this section, we describe the major strategies our participants adopted to solve the complex VAST challenge.

5.1 Starting from a First Clue

The introductory letter read to the participants explicitly mentioned an initial document; all participants started by finding this initial document. Next, teams chose different strategies to collaborate or divide their work: in nine of the teams, both partners worked together to read the first article: either one member would wait while the other found the article, or both searched for it in parallel. In the other six teams, one team member volunteered to read the document; the other began to search for other terms they believed to be relevant (“conspiracy”), or searched for interesting locations on the provided map (“missile silo sounds interesting”).

5.2 Iterative Searching and Reading

After working on the first document, teams engaged in an iterative process of searching for and reading documents they had found. The task required participants to connect facts from a number of different documents; this is reflected in the large number of documents that participants searched for, found, and read. Teams issued on average 50 searches with 42 unique search terms. An average of 90 of the 240 documents were then selected from these searches. 70% of these documents were opened and read by one or both team members (63 documents on average). Interestingly, 58 documents on average were read more than once. We observed that participants revisited previous documents after they or their collaborator had an insight about their work so far, or found new information in other documents. The remaining 30% of all documents pulled out of the list were never opened but left on the tabletop. We suspect that these documents were kept as bookmarks or items to be checked later, and then were later determined to be unnecessary.

5.3 Information Sharing and Collaboration

As participants worked on the dataset, they adopted different collaboration styles. At times, they would work on the same problem, even reading a single document together; at other times, they would separate to work on different problems. These different collaboration styles allowed them to flexibly investigate temporary hypotheses, test ideas, and continue to build a narrative. In order to more closely understand which information and views of the data space participants were sharing and collaborating on at a given moment, we coded their interactions based on their data views and personal interactions. We identified the following eight different collaboration styles and activities that participants adopted (see Figure 4 for illustrations on each collaboration style):

DISC: Active **discussion** about the data or task. Limited system interaction (e.g., pointing to items or scrolling in documents).

- VE:** **View engaged.** One person is actively working; the other watches and engages in conversation and comments on the observed activities, but not interacting with the system.
- SV:** **Sharing of the same view** of a document or search result. Participants either look at the same document reader or the same search result list together at the same time.
- SIDV:** **Sharing of the same information** but using **different views** of the data. Participants for example read the same document but using their own copies (views) of the document.
- SSP:** Work is shared to solve the **same specific problem**. Both read **different documents from a shared set**. For example, participants issued a search for “injured driver,” divided the work and each person read one half of the documents.
- SGP:** Work on the **same general problem** but **from different starting points**. For example, both participants search for documents to find information about the injured driver but start from different searches (e.g. “injured driver” and “obituaries”) and consider different sets of documents.
- DP:** Work on **different problems**, and hence different aspects of the task. For example, one person is interested in the injured driver, the other searches for events around the missile silo.
- D:** **Disengaged.** One person is actively working, the other is watching passively or is fully disengaged from the task.

During the video coding we recorded time-stamped event-logs for each of these collaboration categories. We further grouped them into those describing ‘close’ collaboration (the first five codes) and ‘loose’ collaboration (the last three codes). Close collaboration was generally characterized by active sharing of information and discussion of hypotheses. During phases of close collaboration, teams shared these temporary hypotheses, looked at related information, and pursued similar questions. During phases of loose collaboration, explicit verbal sharing of information and evolving hypotheses was less frequent as team members were looking at less related information. Our coding revealed that teams showed high task engagement, with very little time spent disengaged (D) from the task (under 1% of the time on average). Overall, we observed a tendency for groups to share information and hypotheses frequently, with eleven (of 15) groups spending over half of their time in close collaboration. We refer to those eleven groups as closely coupled; the remaining four groups are loosely coupled.

In loosely-coupled teams, participants spent on average 60% of their time working in parallel, and only infrequently interacted with their team members. Loosely coupled teams spent the dominant portion of their time pursuing different searches and ideas, but working on the same general problem (SGP). In SGP, teams were trying to answer a common general question such as “what is the involvement of Boynton laboratories” but starting from different searches. The video coding revealed that 43% of their total task time was spent in this condition. In these groups, team members worked largely in their own part of the workspace with separate search results and documents; they would discuss (DISC) what they had found 7% of the time. The bottom row of Figure 5 shows one loosely-coupled team.

Closely coupled groups spent on average 70% of their time closely interacting with the other team members. They spent the most time working from a shared set of documents (SIDV, taking 24% of their total task time). They were also likely to examine the same views as each other (SV, 13%). Further evidence for close interaction between participants in closely coupled groups is the overall discussion time (DISC) with 11%, and the time spent discussing and analyzing the same views (SV) of the data with 13% of the total task time, compared to just 5% for loosely coupled teams. The top row of Figure 5 shows one closely coupled team. How

much information participants shared and how frequently they connected to others had an influence on how well teams were able to connect the facts and progress through the task. We discuss this in more detail in Section 6.

5.4 Presenting the Solution

Groups were successful in finding the documents required to solve the task. Eleven of the 15 groups found all ten critical documents; the remainder missed one or two. It is worth noting that while the dataset was large, all teams successfully found most of the critical documents in the first half the experiment. While participants had found most of the relevant information by that time, however, they had typically not managed to connect the pieces of information they had into a coherent hypothesis. During the study, the experimenter paused the teams half-way through to find out what they thought was going on. Much to our surprise, members of several of the teams had radically different working hypotheses from their partners: in one dramatic example, one participant had been pursuing a hypothesis centered around a political intrigue, while the other was working on a terrorism question. As a result, the checkpoint turned out to be accidentally invaluable for several of the teams. The checkpoint was sometimes the first time when participants would link subtle details together: one person would express a hypothesis (“it was e-coli poisoning”) and the other would realize that an intuitive leap had been missed (“no, the apples were poisoned!”). After the checkpoint, however, teams began to better find the value of coordinating; almost all of the teams had a unified story by the last stage.

6 THE VALUE OF COLLABORATION

We observed a substantial degree of variance in how much team members chose to work together. Groups spent anywhere from 32% to 92% of their time in close collaboration. This variability manifested in different amounts of information sharing between groups, impacted how Cambiera’s features were used, and also led groups to be more or less successful the task overall. In order to illustrate how different collaboration strategies impacted how well groups performed in the task and how they used Cambiera, we begin by discussing two different teams’ strategies in detail.

6.1 Strategy Examples

Group 2: Working Closely Together

The two participants in this group were close friends and co-workers. They found all ten critical documents, made all eleven connections with no assists, and solved the complete task correctly in 70 minutes. Both team members worked closely coupled for 92% of the time (Figure 5, top). They had a clear work strategy: they searched and browsed the results in parallel but read interesting articles together, and exchanged 13 documents. They were able to rapidly identify connections between facts, and moved through the study very efficiently. Since these participants worked very closely, with frequent communication throughout, they generally had very good awareness of what each other were searching for. Each made sure that their partner read important documents that they had found, so Cambiera’s awareness features were less critical for this group. As a result, the pair used query coloring less to track their own searches, and more as a way of finding documents that looked to be information-rich. In particular, they preferentially read documents that had several stripes (found by several searches), and that had not been read before.

Group 5: Failing to Combine Knowledge:

Participants in Group 5 were friends, co-workers, and experienced puzzle solvers. They had coordinated multi-hundred-person puzzle-solving competitions, and so felt very confident about their ability to solve the mystery. During the 69 minutes of their work,




Figure 4: Teams adopted eight different collaboration styles which influenced how much information was shared among them. In this image, the item in **SV** represents one shared document (view). In **SIDV** both partners look at the same document (same color) but using their own views. In **SSP** team members read different documents from a common search (same shade of color), in **SGP** team members read documents on a common general topic but starting from different searches, and in **DP** both read documents on different topics (different color).




Figure 5: Times two groups spent in close (yellow) and loose (blue) collaboration the time of the task. The top group spent 92% of its time in close collaboration while the group on the bottom just spent 33% of their time closely coupled. White bars indicate times groups spent talking to the experimenter to summarize current hypotheses or ask questions.

they found all critical documents but only found five connections and required assistance three times, which was among the worst performance outcome of the study. Both participants reported that they were accustomed to working separately, trying to figure out a puzzle on their own. As a result, they adopted a loosely coupled work style, working closely coupled only 33% of the time (Figure 5, bottom). After reading the initial document together, they each chose a part of the problem that was of interest to them. As they worked, they would look for documents that might help their current approach, and periodically mentioned their thoughts aloud. They spent substantial portions of the task unaware of each others' work. It was not until the experimenter asked for a status update that they began to realize what information they were missing. The group spent much of the study working separately on different problems (DP), periodically checking in. During the task, both partners issued 78 searches; they later reported that they had gotten lost in the sheer quantity of results they had found. Because they spent most of the time working on different problems, they infrequently found overlapping documents; as a result, they almost never used Cambria's awareness tools.

6.2 The Value of Close Collaboration

After closely studying the behavior of our groups we were interested to find out whether the overall pattern of collaboration was linked to how successfully groups were able to solve the task. We calculated a debrief score as proposed for judging the 2007 VAST contest [17] from the debriefing of our participants. At the end of the study we asked participants to report the current hypotheses they had derived. We recorded correctly reported, incorrectly added, and missing facts to form our score. Figure 6 gives an overview of the percentage of time each group spent closely coupled (vertical axis) and their debriefing score (horizontal axis). Except for two outliers, groups that spent more time in close collaboration generally had a higher debriefing score in our study.

Our study differed from the scoring of the VAST contest, however, in that some groups received assists during the study. Whether a group required an assist is another indication of how well it was able to work on the task. The four groups which required no assists spent over 69% of their time in close collaboration. Loosely

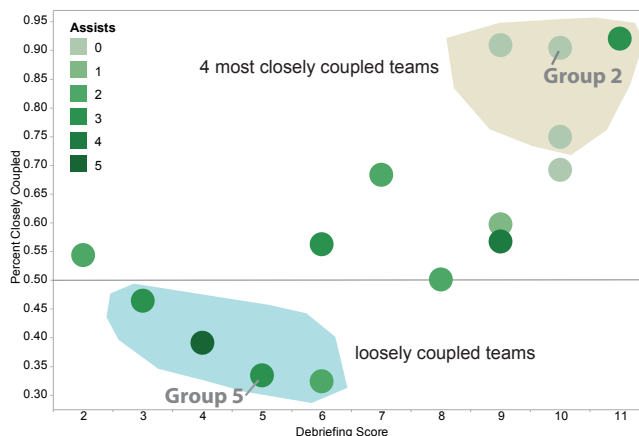


Figure 6: Relation between debriefing score, percentage of time a group spent closely coupled, and number of assists per group.

coupled team required between 2–5 assists. The results for our four most loosely and closely coupled teams are presented in Table 1 and highlighted in Figure 6.

These results suggest that coupling style was linked to teams' success in the task. Pairs that synchronized frequently on their discoveries, in general, did better than those that did not. They connected more facts and required fewer assists. In the task, many of the connections between documents and ideas were subtle and required a leap of insight. For example, participants had to realize that what seemed to be e.coli symptoms in one document might well be a poisoning in another document. Pairs who worked on the same problem ("let's see if we can find other discussions of the poison") were better able to make these connections, while pairs who worked completely independently ("I'll work on the corruption story; you work on the murder") appeared to be less successful. For example, Group 5 (discussed in the previous section) is one example of a group that did not work closely together, and did not perform well

Table 1: Debriefing score and assists for the four loosely coupled teams (top) and the four most closely coupled teams (bottom).

Group	Debriefing Score	Assists	Closely Coupled
7	6	2	32%
5	5	3	33%
8	4	5	39%
3	3	3	46%
12	10	0	74%
2	10	0	90%
15	9	0	90%
1	11	3	92%

on the task overall (see Figure 6). The summaries that they gave each other infrequently included the critical facts that they needed to make connections. Similarly, working too separately was a detriment to other pairs.

Overall, teams that were more successful—that is, reported more correct and fewer incorrect facts, and got fewer hints—were ones that spent more time closely coupled. Note that this correlation does not suggest causation: that is, it is unclear whether success on the task leads to coupling or coupling leads to success. It is likely to be a combination of both: in many teams, we observed pairs of participants working independently until one of them found something that seemed promising. Frequently, finding a critical document during separate work led pairs to switch to a more closely coupled style to discuss intermediate results, to read the critical document or other related documents together, to interpret the facts found up to that point, or to offer/ask for help. This was echoed during the debriefing session where eight of the fifteen pairs explicitly stated that they tended to work in a form of divide-and-conquer style. They would first go off and do individual work and then come together when they had found a partial solution or hypothesis to report or wanted to get the other person involved for validation.

6.3 Suitability of the Tabletop as a Collaborative Visual Analytics Workspace

During the study, the tabletop was the collaborative center of users’ engagement in the task. Cambiera was designed to take advantage of the tabletop setting through specific features that allowed participants to share and organize information freely similar to what would be possible in a physical setting. Two key benefits arose from using tabletop that would not have been as accessible in either a multiple-display or a shoulder-to-shoulder configuration: artefact-centered information sharing and face-to-face interactions. Both sharing and face-to-face work enhanced collaboration: partners were able to point to and manipulate documents and searches that they felt were relevant, and point to documents that their partners could see. During discussion (DISC) periods, team members did not need to turn away from the table in order to talk face-to-face but instead discussions stayed focused on the shared information artefacts, while being able to easily pull in and compare related information. In the debriefing session, one team member said: *“sometimes [my partner] and I fight about who gets to use the computer so it was nice to share that. One person standing and another sitting that’s never a collaborative act, here we’re assembling a whole thing together.”* Others stated that working face-to-face helped them, *“we kept talking the whole time and we assigned things to each other.”*

Several team members volunteered that they particularly liked the direct-touch nature of the tabletop for sharing electronic documents. In the debriefing session, 11 of the 15 pairs highlighted the ease of sharing documents as one of their favorite and most useful features of Cambiera. Five pairs expressed that they particularly en-

joyed working in a face-to-face setting. Last, the tabletop seemed to allow people to think about the task in a different way: *“I liked the direct touch: it’s more human, less technical.”* While some of this enthusiasm can be attributed to the novelty of touch input, all teams approached the table without hesitation and felt encouraged to interact with the system.

The face-to-face tabletop setting and the ability for both partners to synchronously search and read was a strong benefit for the complex problem solving context. We observed several instances in which the collaborative setting contributed to participants making progress in the task. Team members shared and reflected on each others’ ideas, cross-validated hypotheses by sharing documents and results, and engaged in repeated discussion about current hypotheses and strategies on how to proceed further. We observed all pairs talking to each other, gesturing toward searches and documents on the tabletop, and generally utilizing the shared reference point of having a tabletop in front of them. In the questionnaire, 23 of 24 participants “agreed” or “strongly agreed” that they knew what their partners were working on, suggesting that the tabletop provided a strong context for mutual awareness.

We note that some teams felt that the Microsoft Surface tabletop, in its current form, was a constraint. At 2’x3’, and 1024x768 pixels, team members often felt cramped, wanting a higher-resolution and physically larger display for document reading. The lack of a physical keyboard meant that typing was more cumbersome and error-prone than on a physical keyboard. While many of the tabletops on the market suffer from similar difficulties, some projects are beginning to explore larger-scale work surfaces [8] and integrated keyboards.

7 IMPLICATIONS FOR DESIGN

In summary, our study found that working face-to-face around the table was a successful way for pairs to solve the complex problem. The setting allowed participants to approach the problem quickly and effectively. All participants immediately immersed themselves in the task and made use of the various features Cambiera offered. One of the surprising findings of our study was that task success and time spent working closely together were highly correlated. This was not a challenge of information finding: every team found most or all of the critical documents; the only question was whether they were able to draw the connections between them.

Based on our findings, we draw some design implications for improving Cambiera as well as other co-located collaborative problem-solving tools.

7.1 Design for Transient Behavior

The complex problem required collaborators to react to emerging information and hypotheses and required the team to frequently reassess and change their current work strategies. As there was no obvious best way to solve the task, teams often started out with different strategies. The support of a wide variety of work styles and collaboration strategies is a challenge for the design of collaborative problem solving tools. It is insufficient to support just one strategy, since different team members employ different strategies, often in parallel, and pairs frequently shift strategies depending on the current stage of the task. Cambiera, for example, offered flexibility through free workspace organization and repositioning of searches and documents. Team members used Cambiera to engage in different analysis activities in parallel: to search, read, and extract information at any given time without influencing others’ work. By allowing team members to create their own copies of documents and search results as needed, and by not forcing participants to explicitly share resources, these flexible work behaviors were able to emerge. We recommend that collaborative systems continue to allow teams to smoothly choose which collaboration style best fits their task requirements and work styles.

Cambiera, however, did not offer sophisticated storage or sharing mechanisms, but participants did request features for relating data to each other and sharing partially-assembled conclusions. The challenge in offering these mechanisms is to design them in a way that they can easily and fluidly be changed, re-appropriated, or morphed depending on task progress and current questions that the team is trying to answer. For example, we envision interactive mechanisms that can easily rearrange facts based on different semantic arrangements as well as team members' preferred placements. Other collaboration tools for complex tasks could similarly benefit from flexible methods to arrange and share information, even outside of an intelligence analysis task.

7.2 Encourage Closely Coupled Work

Pairs were able to take advantage of the unique affordances of the tabletop—the face-to-face interaction and direct manipulation of objects—to organize their thoughts, share ideas, and work through the problem. Pairs that worked for too long in parallel were less successful in connecting the facts they found, and so had trouble solving the challenge. The awareness features built into Cambiera were meant to bring teams closer together by encouraging conversation about documents that both team members had read or found. While participants found this subtle awareness information useful and valuable, we conclude that stronger indicators would be more beneficial. These could make more obvious visual connections between common information that team members are reading, that are placed in close proximity, or that have been stored in the workspace. Collaborative annotations and notepads may also have encouraged the loosely coupled teams to re-connect.

8 CONCLUSIONS

We presented a detailed exploratory study of a complex collaborative problem solving activity around a digital tabletop display. We studied Cambiera, a tool for collaborative analysis work with text document collections. The contributions of this study were a set of findings on our digital tabletop setting as a context for co-located problem solving. We explored its suitability, identified eight collaboration styles which pairs adopted while solving the problem collaboratively, and described how collaboration impacted their success in the task. In particular, our study showed that Cambiera—in the digital tabletop setting—was a successful work context for complex problem solving. Our task required teams to constantly react to new information, to re-interpret what they had found, and to re-assess their strategies. Hence, participants worked together in a variety of work styles, supported by Cambiera's flexible collaborative search, organization, and sharing mechanisms. We found that teams that connected most often about their individual findings, and worked closely together throughout, were more successful at the task and required fewest assists. Based on these observations, we offered recommendations for features that could be used to improve co-located problem-solving tools more generally. In particular, support for teams to make ad-hoc changes to all aspects of their current work strategies as well as features that encourage them to share information and connect with each other frequently, are worth considering.

ACKNOWLEDGEMENTS

We thank our participants for taking part in the study, our co-workers at Microsoft for participating in the pilot studies, Sheelagh Carpendale for valuable feedback on the paper, and Chris North for introducing us to the VAST challenge data and tasks.

REFERENCES

- [1] C. Andrews, A. Endert, and C. North. Space to Think: Large High-Resolution Displays for Sensemaking. In *Proc. of CHI*, pages 55–64, New York, USA, 2010. ACM.
- [2] A. D. Balakrishnan, S. R. Fussell, and S. Kiesler. Do visualizations improve synchronous remote collaboration? In *Proc. of CHI*, pages 1227–1236, New York, USA, 2008. ACM.
- [3] A. D. Balakrishnan, S. R. Fussell, S. Kiesler, and A. Kittur. Pitfalls of information access with visualizations in remote collaborative analysis. In *Proc. of CSCW*, pages 411–420, New York, USA, 2010. ACM.
- [4] S. E. Brennan, K. Mueller, G. Zelinsky, I. Ramakrishnan, D. S. Warren, and A. Kaufman. Towards a Multi-Analyst, Collaborative Framework for Visual Analytics. In *Proc. of VAST*, pages 129–136, Los Alamitos, USA, 2006. IEEE Comp. Society.
- [5] G. Chin, Jr., O. A. Kuchar, and K. E. Wolf. Exploring the analytical processes of intelligence analysts. In *Proc. of CHI*, pages 11–20, New York, USA, 2009. ACM.
- [6] C. Forlines and C. Shen. DTLens: Multi-user Tabletop Spatial Data Exploration. In *Proc. of UIST*, pages 119–122, New York, USA, 2005. ACM Press.
- [7] G. Grinstein, T. O'Connell, S. Laskowski, C. Plaisant, J. Scholtz, and M. Whiting. VAST 2006 Contest—A Tale of Alderwood. In *Proc. of VAST*, pages 215–216, Los Alamitos, USA, 2006. IEEE Comp. Society.
- [8] B. Hartmann, M. R. Morris, H. Benko, and A. D. Wilson. Augmenting Interactive Tables with Mice & Keyboards. In *Proc. of UIST*, pages 149–152, New York, USA, 2009. ACM.
- [9] P. Isenberg and S. Carpendale. Interactive Tree Comparison for Co-located Collaborative Information Visualization. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1232–1239, Nov./Dec. 2007.
- [10] P. Isenberg and D. Fisher. Collaborative Brushing and Linking for Co-located Visual Analytics of Document Collections. *Computer Graphics Forum*, 28(3):1031–1038, Juni 2009.
- [11] P. Isenberg, M. Sedlmair, D. Baur, T. Isenberg, and A. Butz. Proceedings of the VisWeek workshop on Collaborative Visualization on Interactive Surfaces—CoVIS 2009. Technical Report LMU-MI-2010-2, Ludwig Maximilians University Munich, Munich, April 2010.
- [12] P. Isenberg, A. Tang, and S. Carpendale. An Exploratory Study of Visual Information Analysis. In *Proc. of CHI*, pages 1217–1226, New York, NY, USA, 2008. ACM Press.
- [13] P. E. Keel. Collaborative Visual Analytics: Inferring from the Spatial Organization and Collaborative Use of Information. In *Proc. of VAST*, pages 137–144, Los Alamitos, USA, 2006. IEEE Comp. Society.
- [14] B. Mirel. *Interaction Design for Complex Problem Solving*. Morgan Kaufmann, San Francisco, 2004.
- [15] M. R. Morris, J. Lombardo, and D. Wigdor. WeSearch: Supporting Collaborative Search and Sensemaking on a Tabletop Display. In *Proc. of CSCW*, pages 401–410, New York, USA, 2010. ACM.
- [16] S. A. Paul and M. R. Morris. Cosense: Enhancing sensemaking for collaborative web search. In *Proc. of CHI*, pages 1771–1780, New York, USA, 2009. ACM.
- [17] C. Plaisant, G. Grinstein, J. Scholtz, M. Whiting, T. O'Connell, S. Laskowski, L. Chien, A. Tat, W. Wright, C. Görg, Z. Liu, N. Parekh, K. Singhal, and J. Stasko. Evaluating Visual Analytics at the 2007 VAST Symposium Contest. *IEEE Computer Graphics and Applications*, 28:12–21, 2008.
- [18] A. Robinson. Collaborative synthesis of visual analytic results. In *Proc. of VAST*, pages 67–74, Los Alamitos, CA, USA, 2008.
- [19] J. Stasko, C. Görg, and Z. Liu. Jigsaw: Supporting Investigative Analysis Through Interactive Visualization. *Information Visualization*, 7:118–132, 2008.
- [20] A. Tang, M. Tory, B. Po, P. Neumann, and S. Carpendale. Collaborative Coupling over Tabletop Displays. In *Proc. of CHI*, pages 1181–1290, New York, 2006. ACM Press.
- [21] F. B. Viégas, M. Wattenberg, F. van Ham, J. Kriss, and M. McKeon. Many Eyes: A Site for Visualization at Internet Scale. *IEEE Transactions on Visualization and Computer Graphics*, 12(5):1121–1128, Nov./Dec. 2007.
- [22] D. Wigdor, H. Jiang, C. Forlines, M. Borkin, and C. Shen. WeSpace: The Design Development and Deployment of a Walk-up and Share Multi-Surface Visual Collaboration System. In *Proc. of CHI*, pages 1237–1246, New York, USA, 2009. ACM Press.