

Note-taking in co-located collaborative visual analytics: Analysis of an observational study

Information Visualization
11(3) 190–204
© The Author(s) 2012
Reprints and permission:
sagepub.co.uk/journalsPermissions.nav
DOI: 10.1177/1473871611433713
ivi.sagepub.com



Narges Mahyar, Ali Sarvghad and Melanie Tory

Abstract

In an observational study, we noticed that record-keeping plays a critical role in the overall process of collaborative visual data analysis. Record-keeping involves recording material for later use, ranging from data about the visual analysis processes and visualization states to notes and annotations that externalize user insights, findings, and hypotheses. In our study, co-located teams worked on collaborative visual analytics tasks using large interactive wall and tabletop displays. Part of our findings is a collaborative data analysis framework that encompasses record-keeping as one of the main activities. In this paper, our primary focus is on note-taking activity. Based on our observations, we characterize notes according to their content, scope, and usage, and describe how they fit into a process of collaborative data analysis. We then discuss suggestions to improve the design of note-taking functionality for co-located collaborative visual analytics tools.

Keywords

Note-taking, collaboration, interactive surfaces, history, provenance

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. [1] argue that there is a critical need to support insight externalization through mechanisms such as taking notes, saving views of data, and annotating views. However, support for these activities in visualization tools is currently limited. In this paper, we discuss the results of a user study that emphasize the significance of record-keeping activities during collaborative visual analytics on interactive surfaces. Our focus is on co-located work by small groups of known collaborators, as illustrated in Figure 1.

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 [1,2]. 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 [2,3]. 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

Department of Computer Science, University of Victoria, Canada

Corresponding author:

Narges Mahyar, Department of Computer Science, University of Victoria, Room 504, 3800 Finnerty Road, Victoria, BC, Canada V8P 5C2.

Email: nmahyar@cs.uvic.ca



Figure 1. Examples of note-taking activities during our observational study. Sometimes note-takers are disconnected from group activities.

more closely [4]. However, co-located collaborative visualization is a relatively new area of research and still considerably underexplored. There are many special challenges and requirements for the design of co-located collaborative visual analytics tools and interaction techniques [3]. Different bodies of work have addressed some of these requirements, but only a few co-located collaborative visualization tools have been developed (e.g. [3,5–9]). Other systems implemented for collaborative visualization, such as [10–12], have focused on distributed work, which has different requirements. None of these tools so far has focused on record-keeping activities during collaboration.

Better support for record-keeping activities emerged as a critical need during our observational study, prompting us to analyze note-taking (i.e. using a pen to write something on a piece of paper) and other record-keeping activities in depth in order to establish design requirements. This paper presents the study and results, and offers suggestions about how to improve the design of co-located collaboration tools. The results of our study were first published at the VAST 2010 conference [13]. In this extended version, we elaborate on our results, discussion, and relationships to previous literature.

We use the term *visual representation* to refer to artifacts that display data, *visualization* to refer to the process of creating and editing visual representations, and *visual analytics* to refer to the larger process of using visual representations and other sources of information to form insight and make decisions. Our findings indicate that record-keeping is a pivotal activity that is carried out throughout a data analysis session. We propose a categorization of notes based on their content, scope, and usage, and discuss how record-keeping fits into the visual analytics process. We then discuss related work, comparing our categorization and findings with others reported in the literature. Finally, we discuss potential ways in which note-taking could be integrated into collaborative visual analytics tools and present some design suggestions.

Method

Our laboratory study examined the process of co-located visual analytics in a business context. The study was exploratory in nature rather than designed to test a specific hypothesis. Our goal was to better understand collaborative activities and challenges that might suggest improvements for collaborative analytics tools. 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 wall and tabletop displays.

Task

Each group completed two tasks, both using a sample e-fashion dataset from Explorer [14]. 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.

We received advice from business professionals and faculty members in designing our tasks. Because our participants did not own the data and were not familiar with it, we were concerned about their engagement. Hence, we decided to first familiarize them with the data and the interface by using focused questions in Task 1. We then designed Task 2 to involve competition among the members of a group, which we hoped would engage the participants in the analysis process. Engagement was

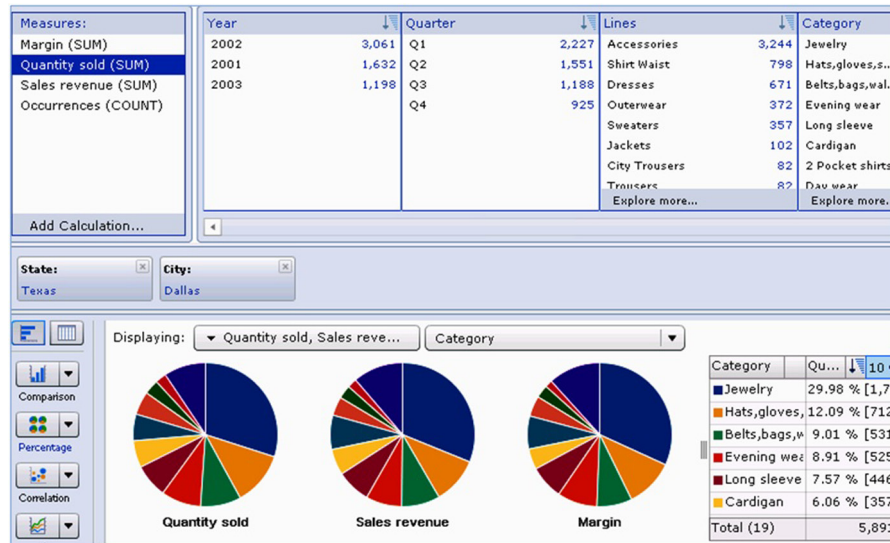


Figure 2. Partial screenshot of Explorer, depicting a comparison chart that visualizes margin, quantity sold, and sales revenue over category, filtered according to state (Texas) and city (Dallas).

high, but the competitive nature of Task 2 did have an impact on the analysis process, as described later.

Participants

Twenty-seven student participants took part in our 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, we 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 users had experience with some kind of data analysis software such as Microsoft Excel.

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 \times 34.4 inches (70 inch diagonal) and had four high-definition projectors to create a total resolution of 3840 \times 2160.

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

We used “Explorer” [14] (Figure 2) as our visual analytics tool. 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. We therefore expected some problems when using it collaboratively. We 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.

Procedure

We began with a 10–15 min introduction to Explorer. Participants then spent approximately 30 min on Task 1 and 40 min on Task 2. We offered an optional 5 min break between the two tasks. After Task 2, participants spent approximately 10 min summarizing and writing down their results. We asked participants to create a report of their results at the end of Task 2 to justify their decisions. Following the computer-based tasks, we 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, though participants were encouraged to use the styli for interaction.

Data analysis

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

Two authors coded and analyzed the data together using a structured qualitative data analysis process. We re-coded and re-analyzed data iteratively to refine our characterizations of collaborative processes, activities, and notes. In the first pass, 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. Our main observation in the first pass 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.

Findings

As our study was exploratory, we did not have any particular hypotheses. However, we had anticipated problems such as incorrect software orientation and inadequate awareness of other users' work. We predicted these challenges to be the most important barriers and planned to concentrate on them, but based on our observations we found note-taking a more interesting and yet less discussed obstacle to investigate.

Participants' collaboration and use of software

We 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 [15] 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 we cannot be sure of the magnitude and significance of these effects. As the software layout dictated positioning, we 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 [11], 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 we 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.

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

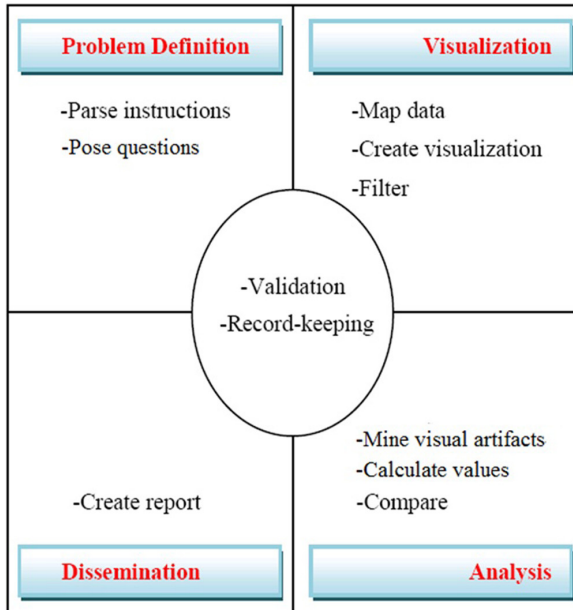


Figure 3. Activities we observed within the problem definition, visualization, analysis, and dissemination phases of collaborative visual analytics. Activities in the middle circle are common in all four phases.

Phases and activities

We noticed a similar analytics process among all of the groups that we observed. We characterized groups' actions at two levels: high-level phases and low-level activities, as shown in Figure 3. As this characterization is grounded by the particular data that we collected, we cannot generalize the phases and activities to other data analysis situations. However, similarity to other frameworks [16–19] suggests that many aspects of this process probably occur outside the context of our study.

We identified four high-level phases: **problem definition, visualization, analysis, and dissemination**. Figure 3 shows that there are activities common to all phases such as record-keeping and validation, and activities unique to each phase. We explain each phase and their exclusive activities below. Common activities are explained separately.

Our findings confirmed what Isenberg et al. [16] stated about the non-linear temporal order of activities. In our study, we 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.

Phase II: visualization. We use the term visualization to describe a group of activities resulting in a visual artifact (i.e. a chart). The visualization phase comprises a number of essential activities, specifically:

- mapping variables;
- filtering underlying data; and
- 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 our 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; and
- 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

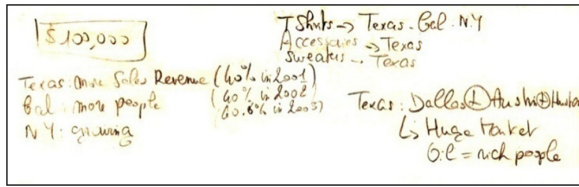


Figure 4. Note taken by a participant consisting of values such as \$100,000, calculated numbers such as 60%, and externally acquired information such as “Texas: Dallas + Austin + Houston → Huge market.”

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 Internet, etc. Figure 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.

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. We 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 capturing analytic activities and visualization states to notes and annotations. In our 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.

Record-keeping strategies

Table 1 reports quantitative data derived from our 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 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.

Our 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. We focus only on the two extreme approaches.

We 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 our 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

Table 1. Number of note-taking and chart-saving actions by each group. Shaded groups relied heavily on saved charts for analysis.

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

N: no; Y: yes.

for 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 1). Groups who saved charts for the second purpose saved a larger number of charts. We 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. We noticed that the overall amount of note content taken by participants in Group 2 was less than that for the other groups. Based on our analysis of their activities, we 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.

Characterization of note-taking activities

Notes' content. Based on our analysis of the notes taken by participants, we 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 our 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. We 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, we 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 ↑ to indicate the increase of a value such as revenue.

California 29.79% (28.9%)	New York 20.98% (30.6)	Texas 27.64%
\$ 1.121 million (8M)	\$ 1.189 million	\$ 1.566 million
17769 (OS)	19109	25,000
8935 - sweat + shirts	7736 Sweat T's (↔)	11229
2176 - Accessories	3279 Accessories (↔)	3612
1967 - shirt waist	2165 Sweaters (↔)	3211
\$ 335,000 Sweat t's	\$ 558,000 Sweat t's (↔)	\$ 804,000
\$ 113,000 Acces	\$ 182,000 Acces (↔)	\$ 187,000
\$ 115,000 shirt waist	\$ 145,000 Sweaters (↔)	\$ 161,000
location \$ 29,000	\$ 31,000	\$ 40,000

Figure 5. A note that was taken for group use. It is nicely structured and comprises information for all three rivaling participants.

Notes' scope. Based on the way that notes were shared, we divide them into group and personal notes. We 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.

We noticed that the nature of the problem influenced the scope of notes. During Task 1 (which had a cooperative 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 our 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 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.

In contrast, Figure 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 6 (a)). Similarly, other participants also noted interesting findings for their own states (Figure 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 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.

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

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.

① Sweet T-shirt
 Texas 42 50% T-shirt
 New York 30
 California 28

(a)

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

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.

Calo. Revenue 2.2 million 2.8 1.7
 67
 28 | → 40%
 20 | → 20%
 19 | → 28% 26%

(b)

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

4.2 mil sales revenue 2003

(c)

Figure 6. Three personal notes taken by participants in one group. This group did not have a designated note-taker.

high Q3 sales				
<u>New York</u>	<u>Q1</u>	<u>Q2</u>	<u>Q3</u>	<u>Q4</u>
o1	555	480	257	371
	684	613	500	886
	747	856	914	637
<u>Texas</u>	<u>Q1</u>	<u>Q2</u>	<u>Q3</u>	<u>Q4</u>
o1	559	615	329	496
o2	1014	796	754	1438
o3	1102	1088	1032	962
<u>Cali</u>	<u>Q1</u>	<u>Q2</u>	<u>Q3</u>	<u>Q4</u>
o1	579	441	394	349
	650	549	760	842
	729	787	775	698

Figure 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.

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 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.

Awareness with respect to note taking

We noticed that the manual note-taking process impacted awareness among group members. This can be seen in Figures 1 and 8, where note-takers are disconnected from the other team members. In both cases, the note-taker (female) 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 5 was assigned the role of note-taker. He was sitting most of the time observing




Figure 8. A seated note-taker of a group working on the wall display had her view obscured by two other participants.

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.

Wall display versus tabletop display

Table 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 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 [4], our 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, we 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 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.

Related work

In this section, we first present work that addresses the collaborative visual analytics process and requirements. We then discuss related work that specifically highlights the importance of history mechanisms and note-taking during the course of analysis.

Collaborative visual analytics process

While substantial research has been devoted to computer-supported cooperative work in general, collaborative visual analytics is still underexplored and has unique challenges. Researchers have identified the need for very flexible tools [16,20,21], including flexibility to change ordering of activities, work styles (from closely coupled to independent), role assignments, and workspace organization [7,16].

Several studies have examined how users analyze data to characterize the processes and activities involved (e.g. [22]). More relevant to our work are studies that consider analytic processes of groups by using software supporting collaborative work [18,19] or by using paper-based tasks [11,20]. Findings of previous studies, regardless of whether the tasks were paper based or software based, resulted in similar lists of activities. For instance, Mark and Kobsa [18] identified processes of parsing the question, mapping variables, finding or validating a visual representation, and validating the entire analytic process. Isenberg et al. [16] identified processes of browsing, parsing, discussing collaboration style, establishing task strategy, clarification, selecting, operating, and validating. These lists of activities bear strong resemblance to our own characterization. For instance, our first phase, “problem definition,” has been identified in previous work as “parsing” [16,18] or “problem interpretation” [19].

In contrast to previous work, our 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. We believe this two-level structure provides a useful way to think about the analysis process. We also highlight record keeping as a critical activity during all phases. Although record keeping has been previously mentioned as a relevant action [17,23,24], its importance may have been underrepresented in previous frameworks describing collaborative analytics processes.

Record keeping

Here we describe relevant work related to saving charts/system states and taking notes, as these were the two record-keeping strategies that we observed.

History of states. In visual analytics tools, record keeping 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 [4,10,25–27]. According to Shneiderman [24], 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. [10] noted, history tools can also be used to create a report or presentation after analysis is complete. Isenberg and Carpendale [3] stated that while data analysis histories are necessary for individuals, they might be of even higher importance in collaborative settings. 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 [28]. In addition, data exploration history can be used at a later time to discuss or share interesting findings [3].

We further emphasize that the vast majority of history/provenance tools have focused on single-user systems. Although previous work has postulated that history tools may be even more important for collaborative work [2,29], little guidance is available to help build such tools effectively. Extending history mechanisms to represent activities of multiple co-located users is non-trivial owing to issues of awareness, disruption, organization, and so on. In a previous workshop paper [30], we hypothesized how history tools might need to change to support multiple users.

Note-taking and annotation. 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 [31]. 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 [32]. Furthermore, note-taking seems to assist complex tasks such as problem-solving and decision-making by reducing the load on working memory [32]. It has also been observed that taking notes keeps students engaged and improves the learning process [32,33].

The importance of note-taking and annotation in visual analytics has also been mentioned [3,10,34]. Heer et al. [10] stated that annotations and notes are important for supporting discussions around visualizations in

distributed collaborative visual analytics. Kadivar et al. [35], 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 [36]. Furthermore, notes help to create a link between the system and an analyst's cognitive processes [37]. Lipford et al. [1] stated that externalization improves recall at a later time, which helps analysts leverage their previous findings. Externalization refers to recording of insights in the form of notes, annotations, and bookmarks.

Insight externalization has been implemented in some research tools. Sense.us (see [23]) allows users to collaboratively analyze data and add annotations on top of visualizations and write down their findings in notes attached to the visualization. Aruvi (see [37]) enables users to take notes and link related notes to each other and the visualization. Harvest (see [11]) has a note-taking mechanism that automatically recommends the notes that are most related to an analyst's current line of inquiry. Collaborative Annotation on Visualization (see [38]) allows analysts to remotely collaborate and add annotations on top of their visualizations. Note that the purpose of note-taking can be quite different in distributed collaboration than in co-located 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.

Research by Shrinivasan and van Wijk [37] suggests 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. Shrinivasan and van Wijk [37] take annotation a step further by automatically recommending related notes based on the current analysis context. Our findings highlight note-taking as a pivotal activity during the course of analysis, emphasizing the importance of including such provenance tools in visual analytics systems.

Discussion

Our observations show that both taking notes and saving charts were employed by our participants to help them perform analytic tasks. In the next section, we use evidence from our observational study to propose more specific design guidelines and considerations for visual analytics tools, with a focus on note-taking.

A clear need for record-keeping support

Our main result is the importance of recording findings and cues (as notes or saved charts). Although some

previous research [3,23,34] has suggested allowing annotation of visualizations, our 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 we 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 [35,37,39] has demonstrated how this can be done for single users, but work remains to extend this idea to multi-user systems.

Impact of task nature on note taking

Our study further corroborated Heer and Agrawala's [40] proposition that the nature of the task affects both collaboration style and division of workspace. Our 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 our 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 [16,20,41,42]. 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.

Suggestions to support note taking

How to best design effective record-keeping functionality for co-located work is not entirely clear. Here we 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 our 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.

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?” Our 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, we 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 chronologic 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. We 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 chronologic 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, we believe that searching and filtering both types of objects will be important.

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 color 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.

We were interested to note that our 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.

Record keeping for different types of collaboration. We 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.

Note manipulation and management. Note management and manipulation is also important to consider. Highlighting, grouping, and summarizing notes will increase their usability [33,37,43]. 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 [32]. We also observed in our user study that most of the group notes were naturally taken in tabular format (Figure 5), which facilitated comparison.

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. We do not believe that there is one perfect solution, but in this section we 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.

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. [34] 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 [44]. The many benefits of tangible interaction are well known [26,45,46]. 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.

Generalizability

Our results are subject to some caveats. We chose to focus on the business domain, so our users were primarily business students. We suspect that collaborative use of

visualization tools will be similar for other group decision-making tasks, but it is possible that we observed some peculiarities unique to business. Second, we chose to utilize existing visualization software to ensure that users could work with interactive and customizable representations of data. However, our users’ behavior may have been influenced by the available technology, especially their closely coupled work style (owing to the inability to create different charts in parallel) and the tendency of most groups to write notes rather than save charts (owing to the difficulty of saving charts with the tool). Also, although we triangulated our results, additional coders may have led to higher reliability. Finally, we examined a group size of three. Collaborative processes are likely to differ for larger group or pairs as it is well known that size impacts the social dynamics of a group (e.g. [47]).

Conclusion and future work

We characterized phases and activities involved in collaborative visual analytics for co-located groups. We also identified record keeping as a process that is intensively used by data analysts. We characterized notes according to whether they were findings or cues, and whether their scope was for personal or group use. We also described how notes were taken and used within four identified phases of data analysis. These analyses enabled us to offer numerous suggestions of how to better support record-keeping activities within visual analytics tools.

Additional studies are required to answer questions about how exactly note-taking support should be provided in collaborative visualization systems. For instance, it is still unclear how we can best support both individual and group note-taking activities. Another important consideration is the form of input to be used for note-taking in a shared display situation. Further studies should also be conducted in other application domains. For instance, in some disciplines, records of decisions need to be kept for legal purposes, and therefore may need to be more formal and detailed. Further research is needed to investigate how these diverse needs can be best supported within visual analytics tools.

Currently, we are developing a prototype system designed specifically for co-located collaborative visual analytics. In the design of this system we have tried to address general issues of collaborative work such as awareness and territoriality. Moreover, we have designed and implemented modules for tracking user analytic activities as well as note-taking and annotation. Our future user studies will examine our design decisions and the effects of different record-keeping mechanisms on the visual analytics process.

Acknowledgments

We thank our colleagues in VisID (Visual Interaction Design) Lab at the University of Victoria for their suggestions.

Funding

We thank SAP and the Natural Sciences and Research Council of Canada for funding this research.

Declaration of conflicting interests

The authors declare that they do not have any conflicts of interest.

References

- Lipford HR, Stukes F, Dou W, et al. Helping users recall their reasoning process. In: *Visual analytics science and technology (VAST)*. Salt Lake City, UT, 2010, pp.187–194. New York: IEEE.
- Heer J, van Ham F, Carpendale S, et al. Creation and collaboration: engaging new audiences for information visualization. *Inf Visual* 2008; 4950: 92–133.
- Isenberg P and Carpendale S. Interactive tree comparison for co-located collaborative information visualization. *IEEE Trans Visual Comput Graphics* 2007; 13(6): 1232–1239.
- Rogers Y and Lindley S. Collaborating around large interactive displays: which way is best to meet. *Interact Comput* 2004; 16(6): 1133–1152.
- Forlines C, Esenther A, Shen C, et al. Multi-user, multi-display interaction with a single-user, single-display geospatial application. In: *Proceedings of ACM symposium on user interface software and technology*, Montreux, Switzerland, 15 October 2006, pp.273–276. New York: ACM.
- Forlines C and Shen C. DTLens: multi-user tabletop spatial data exploration. In: *Proceedings of ACM symposium on user interface software and technology*, Seattle, WA, USA, 23 October 2005, pp.119–122. New York: ACM.
- Isenberg P and Fisher D. Collaborative brushing and linking for co-located visual analytics of document collections. *Comput Graphics Forum* 2009; 28(3): 1031–1038.
- Tobiasz M, Isenberg P and Carpendale S. Lark: coordinating co-located collaboration with information visualization. *IEEE Trans Visual Comput Graphics* 2009; 15(6): 1065–1072.
- Vernier F, Lesh N and Shen C. Visualization techniques for circular tabletop interfaces. In: *Proceedings of the working conference on advanced visual interfaces*, Trento, Italy, 22 May 2002, pp.257–265. New York: ACM.
- Heer J, Mackinlay J, Stolte C, et al. Graphical histories for visualization: supporting analysis, communication, and evaluation. *IEEE Trans Visual Comput Graphics* 2008; 14(6): 1189–1196.
- Shrinivasan YB, Gotzy D and Lu J. Connecting the dots in visual analysis. In: *IEEE symposium on visual analytics science and technology*, Atlantic City, NJ, USA, 12 Oct 2009, pp.123–130. New York: IEEE.
- Viegas FB, Wattenberg M, van Ham F, et al. ManyEyes: a site for visualization at internet scale. *IEEE Trans Visual Comput Graphics* 2007; 13(6): 1121–1128.
- Mahyar N, Sarvghad A and Tory M. A closer look at note taking in the co-located collaborative visual analytics process. In: *Proceedings of IEEE conference on visual analytics science and technology (VAST 10)*, Salt Lake City, UT, 25 Oct 2010, pp.171–178. New York: IEEE.
- SAP Business Objects. Explorer in the cloud, <https://create.ondemand.com/explorer> (accessed 30 June 2010).
- Tang A, Tory M, Po B, et al. Collaborative coupling over tabletop displays. In: *Proceedings of the conference on human factors in computing systems, CHI 2006*, Montréal, Canada, 22–27 April 2006, pp.1181–1190. New York: ACM.
- Isenberg P, Tang A and Carpendale S. An exploratory study of co-located collaborative visual analytics around a tabletop display. In: *Proceedings of the conference on human factors in computing systems, CHI 2008*, Florence, Italy, 5–10 April 2008, pp.1217–1226. New York: ACM.
- Card S, Mackinlay JD and Shneiderman B. *Readings in information visualization: using vision to think*. San Francisco, CA: Morgan Kauffman, 1999.
- Mark G and Kobsa A. The effects of collaboration and system transparency on CIVE usage: an empirical study and model. *Presence Teleop Virt Environ* 2005; 14(1): 60–80.
- Park KS, Kapoor A and Leigh J. Lessons learned from employing multiple perspectives in a collaborative virtual environment for visualizing scientific data. In: *Proceedings of conference on collaborative virtual environments*, San Francisco, California, USA, 10 September 2000, pp.73–82. New York: ACM.
- Robinson AC. Collaborative synthesis of visual analytic results. In: *IEEE symposium on visual analytics science and technology*, Columbus, Ohio, USA, 19 October 2008, pp.67–74. New York: IEEE.
- Scott SD, Grant KD and Mandryk RL. System guidelines for co-located, collaborative work on a tabletop display. In: *Proceedings of conference on computer supported cooperative work*, Helsinki, Finland, 14 September 2003, pp.159–178. Norwell, MA: Kluwer.
- Gotz D and Zhou MX. Characterizing users' visual analytic activity for insight provenance. *Inf Visual* 2009; 8(1): 42–55.
- Heer J, Viégas FB and Wattenberg M. Voyagers and voyeurs: supporting asynchronous collaborative visualization. *Commun ACM* 2007; 52(1): 87–97.
- Shneiderman B. The eyes have it: a task by data type taxonomy for information visualizations. In: *Proceedings of IEEE symposium on visual languages*, Boulder, CO, USA, 3 September 1996, pp.336–343. Los Alamitos, CA: IEEE CS Press.
- Meng C, Yasue M, Imamiya A, et al. Visualizing histories for selective undo and redo. In: *Proceedings of Asia Pacific computer human interaction*, Kangawa, Japan, 15 July 1998, pp.459–464. New York: IEEE.
- O'Hara K and Sellen A. A comparison of reading paper and on-line documents. In: *Proceedings of the human factors in computing systems*, Atlanta, Georgia, USA, 22 March 1997, pp.335–342. New York: ACM.
- Oinn T, Addis M, Ferris J, et al. Taverna: a tool for the composition and enactment of bioinformatics workflows. *Bioinformatics* 2004; 20(17): 3045–3054.
- Gutwin C and Greenberg S. Design for individuals, design for groups: tradeoffs between power and workspace awareness. In: *Proceedings of computer supported cooperative work*

- (CSCW), Seattle, Washington, USA, 14 November 1998, pp.207–216. New York: ACM.
29. Mark G, Carpenter K and Kobsa A. A model of synchronous collaborative information visualization. In: *Proceedings of information visualization*, London, UK, 16 July 2003, pp.373–381. Washington, DC: IEEE Computer Society.
 30. Sarvghad A, Mahyar N and Tory M. History tools for collaborative visualization. In: *Proceedings of workshop on collaborative visualization on interactive surfaces (CoVIS)*, Atlantic City, USA, 22 October 2009, pp.21–23. New York: IEEE.
 31. Hartley J. Notetaking in non-academic settings: a review. *Appl Cognit Psychol* 2002; 16(5): 559–574.
 32. Boch F and Piolat A. Note taking and learning: a summary of research. *WAC J* 2005; 16: 101–113.
 33. Bretzing B. Notetaking and depth of processing. *Contemp Educ Psychol* 1979; 4(2): 145–153.
 34. Isenberg T, Neumann P, Carpendale S, et al. Interactive annotations on large, high-resolution information displays. In: *Proceedings of the IEEE VIS/InfoVis/VAST*, Los Alamitos, CA, 29 October 2006, pp.124–125. New York: IEEE.
 35. Kadivar N, Chen V, Dunsmuir D, et al. Capturing and supporting the analysis process. In: *IEEE Symposium on visual analytics science and technology*, Atlantic City, USA, 22 October 2009, pp.131–138. New York: IEEE.
 36. Marshall CC. Work practices study: analysts and note taking. Unpublished report, <http://www.csd.tamu.edu/~marshall/oswrstudy.pdf> (1990, accessed 30 June 2010).
 37. Shrinivasan YB and van Wijk JJ. Supporting the analytical reasoning process in information visualization. In: *Proceedings of human factors in computing systems (CHI 08)*, Florence, Italy, 5 April 2008, pp.1237–1246. New York: ACM.
 38. Ellis SE and Groth DP. A collaborative annotation system for data visualization. In: *Proceedings of the working conference on advanced visual interfaces*, Gallipoli, Italy, 25 May 2004, pp.411–414. New York: ACM.
 39. Gotz D, Zhou M and Aggarwal V. Interactive visual synthesis of analytic knowledge. In: *2006 IEEE symposium on visual analytics and technology*, Baltimore, MD, USA, 31 October 2006, pp.51–58. New York: IEEE.
 40. Heer J and Agrawala M. Design considerations for collaborative visual analytics. *Inf Visual* 2008; 7(1): 49–62.
 41. Scott SD, Sheelagh M, Carpendale T, et al. Territoriality in collaborative tabletop workspaces. In: *Proceedings of conference on computer supported cooperative work*, Chicago, Illinois, USA, 16 November 2004, pp.294–303. New York: ACM.
 42. Jankun-Kelly TJ, Ma K-L and Gertz M. A model and framework for visualization exploration. *IEEE Trans Visual Comput Graphics* 2007; 13(2): 357–369.
 43. Kiewra KA, Benton SL, Kim S, et al. Effects of note-taking format and study technique on recall and relational performance. *Contemp Educ Psychol* 1995; 20: 172–187.
 44. Maquil V, Psik T and Wagner I. The ColorTable: a design story. In: *Proceedings of tangible and embedded interaction*, Bonn, Germany, 18 February 2008, pp.97–104. New York: ACM.
 45. Sellen AJ and Harper RH. *The myth of the paperless office*. Cambridge, MA: MIT Press, 2003.
 46. Terrenghi L, Kirk D, Sellen A, et al. Affordances for manipulation of physical versus digital media on interactive surfaces. In: *Proceedings of human factors in computing systems*, San Jose, California, USA, 28 April 2007, pp.1157–1166. New York: ACM.
 47. Simmel G. The number of members as determining the sociological form of the group. *Am J Sociol* 1902; 8: 1–46.