

Observations of Record-Keeping in Co-located Collaborative Analysis

Narges Mahyar
University of Victoria
nmahyar@cs.uvic.ca

Ali Sarvghad
University of Victoria
asarv@cs.uvic.ca

Melanie Tory
University of Victoria
mtory@cs.uvic.ca

Tyler Weeres
University of Victoria
tweeres04@gmail.com

Abstract

Record-keeping is known to facilitate visual data analysis in single user and asynchronous collaborative settings. We Implemented CoSpaces, a tool for collaborative visual data analysis with a record-keeping mechanism that enables tracking of analysis history. Then we conducted an observational study with ten pairs analyzing a sales dataset, to study how collaborators use visual record-keeping during co-located work on a tabletop. We report actions on visual record-keeping and inferred key user intentions for each action. Actions and intentions varied depending on the analytical phase and collaboration style. Based on our findings, we suggest providing various views of recorded material, showing manually saved rather than automatically saved items by default, enabling people to review collaborators' work unobtrusively and automatically recommending items related to a user's analytical task.

1. Introduction

Working collaboratively can facilitate analysis of complex datasets and may also improve the quality of the work [1]. As such, there is recent interest in designing visualization tools to support collaboration.

Record-keeping is one of the important aspects of collaborative work. Visual Record-Keeping (VRK) refers to the process of capturing and visually representing a data analysis history and a user's externalizations (notes, annotations, etc.). In Visual Analytics (VA), a record-keeping repository typically consists of recorded visualization snapshots, system states, notes, and annotations. For single users, record-keeping helps to generate insights and synthesize knowledge [2],[3]. While data analysis histories might be more important for collaborative tasks [4], the role and value of VRK in collaborative context is less well studied and prior research has been mainly focused on asynchronous distributed VA. In this setting, VRK can facilitate building of common ground, sharing of visual artefacts and findings and facilitate offline discussions [3],[4]. However, intrinsic attributes of synchronous

collaboration, such as transitions between collaboration styles (i.e. loosely-coupled when collaborators work in parallel and closely-coupled when working together) may impose different design requirements for VRK. Benefits and design of VRK in this context are mainly speculative and have not been empirically tested. Moreover, general design guidelines for co-located collaborative VA tools [5],[9],[10],[11], do not include much guidance for record-keeping.

To gain a better understanding of VRK in a synchronous collaborative setting, we designed CoSpaces (Collaborative Workspaces), a prototype tool for co-located collaborative VA on interactive tabletops that incorporates a VRK module. We designed this module based on our previous research [6],[12] and a workshop we organized with business intelligence experts. We conducted a user study with CoSpaces and report users' actions on recorded history items and notes, and their foremost intention/s for those actions. We also describe how the actions relate to analysis phases (information foraging versus discussion) and collaboration styles (loose versus tight). Our focus is on situations in which small groups of people gather face to face to perform visual analysis and discuss their findings. Based on our observations we suggest design considerations that would better tailor record-keeping modules for synchronous collaborative visual data analysis.

2. Background

2.1. VRK in single user and collaborative VA

Many researchers have mentioned the advantages of record-keeping in visualization [8],[9],[10]. Several single user VA tools provide general-purpose undo/redo operations, but this simplest form of record-keeping is inadequate for most complex VA tasks. Heer et al. [2] suggested additionally recording past visualization states, and also suggested (but did not verify) that history improves communication and dissemination of findings. Similarly, note taking tools are important for helping users to record findings and insights.

Several single user VA tools have implemented variations of such record-keeping functionality. For example, Heer et al. [2] integrated a visual record-keeping module into Tableau software (www.tableausoftware.com), that enabled users to visually browse, search, filter and reuse visualizations. Vistrails [11] captures detailed information about scientific workflow, including data, visualizations, and the pipelines used to create the visualizations.

Isenberg and Carpendale [7] stated that while data analysis histories are necessary for individuals, they might be more important for collaborative tasks. Our previous observational study [12] demonstrated that record-keeping is a critical component of the co-located collaborative VA process. However, this study was limited in that it used a system with no built-in record-keeping capabilities. Thus, although the study clearly showed that record-keeping was important, it could not assess whether best practices for single-user history design would extend to co-located collaboration.

Collaborative use of VRK has been mainly investigated in the remote asynchronous context. Heer et al. [3] found that VRK facilitated view sharing, threaded discussions, and social navigation. Similarly, Many Eyes [4] is another web-based tool that enables bookmarking and sharing of views to support discussion.

In the co-located synchronous context, the closest research to ours is Cambiera [5], a tool that tracks each individual's history while they analyze a document corpus. Using colour-coding, past searches and documents are visually represented to increase users' awareness of each other's work. Similarly, we investigate how access to a history of group members' actions influences awareness and discussion; however, we focus on tabular data, which has very different design constraints than a document corpus. In particular, we track users' created charts and notes rather than a record of which documents have been viewed. Somewhat less related is MemTable [13], a smart tabletop surface that captures and visually represents the table contents during meetings including individual participation histories. However, MemTable was designed for more general types of meetings rather than visual data analysis.

2.2. Analytic activities and actions

In a single user context, Gotz et al. [14] identified and categorized various visual analytic behaviours. Their four-tier hierarchy is comprised of tasks, sub-tasks, actions and events. They argue that the action layer carries information regarding users' analytic

intention/s. With a narrower focus, Sarvghad et al. [6] compiled a list of the most probable history operations (browse, search, filter, edit, delete and export).

Isenberg et al. [10] categorized analytic activities in a collaborative context. They derived eight primary visual analysis processes: browse, parse, discuss collaboration style, establish task strategy, clarify, select, operate and validate. Each process contains a number of activities. For instance, while browsing, participants scanned, flipped through and grouped visualizations to gain a better understanding of available information.

In our study, we observed and identified actions on visual record-keeping, and inferred intention/s related to them. Moreover, we report our observations of how these actions were distributed across different analysis phases and collaboration styles.

3. Overview of CoSpaces

CoSpaces was designed for co-located collaborative VA on a large tabletop display. Below we describe primary features of CoSpaces.

3.1. Worksheet

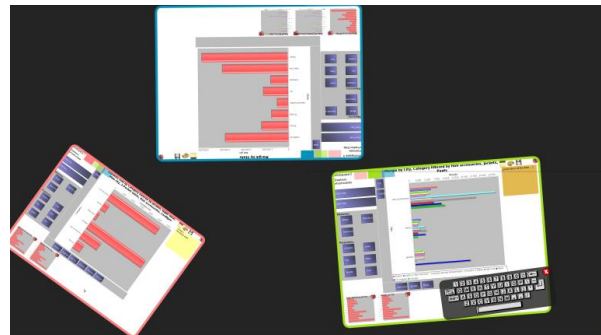


Figure 1: CoSpaces. Dark background is the tabletop surface. There are three open Worksheets.

The CoSpaces interface is composed of Worksheets, as shown in Figure 1. The Worksheet was designed using the principle of “one space, many users”. Its design provides a team with the flexibility to work collectively on one or more Worksheets, or separately and simultaneously on multiple Worksheets. Each Worksheet defines a work territory, either personal or shared. Worksheets therefore enable both individual work territories and shared work territories, as advocated by Scott et al. [15]. Moreover, users may create several Worksheets, perhaps to compare data attributes side-by-side. Personal versus shared Worksheets are identical as far as the system is

concerned; ownership is defined by the way in which they are used. This makes it easy for users to convert a personal space into a shared space or vice versa. Worksheets can also be moved and resized. Each Worksheet's relatively wide border is uniquely coloured with a bright distinctive colour. This enables users to easily distinguish Worksheets from each other. Sections of a Worksheet are shown in Figure 2.

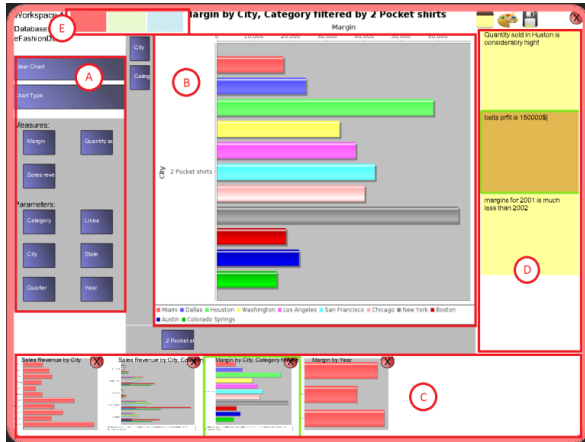


Figure 2: Worksheet Details: Analysis pane (A) for creating and modifying charts, Visualization pane (B), History pane (C), Notes pane (D), and Tabs (E) that provide a portal view to other worksheets.

3.2. Tab Portal Views

CoSpaces uses a tab metaphor to facilitate awareness of other users' activities and share artefacts. Coloured tabs (Figure 2E) are associated with other existing Worksheets. Each tab is colour-coded to match the border colour of the Worksheet that it links to.

Tabs act as portals to view other Worksheets. Tapping on a tab replaces the local worksheet content with a view of another Worksheet. Tapping on the local Worksheet tab switches the view back. When remote viewing, the contents of all panes are changed to reflect the remote information, including the current visualization as well as recorded items in the history and notes panes. The user may browse charts and notes to learn about another user's past analytical activities and interests. To prevent unintentional changes and interruption, a Worksheet's remote view is read-only and navigation between local and remote views is not linked. To share charts, one can select an item in the history pane of a remote view and copy it to the local Worksheet's history pane.

We speculated that remote viewing could have two main uses: (1) to gain awareness of others' work

progress during loosely coupled work, and (2) to review and share findings during closely coupled work.

3.3. Visual Record-Keeping

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

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

A Worksheet automatically captures and saves a copy of the current analysis-state right before a change, made by the user, has been applied. We use a simple heuristic inspired by the chunking rules of Heer et al. [2] to reduce history repository size. An analysis-state is saved only when a change in the current mapping of data takes place. Adding or removing filters will not result in a save.

An analyst can externalize findings, hypotheses and so on using the notes pane. The importance of connecting externalized material to the visual representation of data has been previously recognized [16],[17]. Therefore, we automatically create a link between the current chart and the note.

As part of the analysis-state, we capture a thumbnail picture of the chart. Thumbnails are placed in the history pane in chronological order from oldest to newest (Figure 2C). The pane scrolls as the number of thumbnails grows. Notes are placed in the notes pane in chronological order, matching the chart thumbnails. The notes pane scrolls when the available space is exceeded.

4. Observational Study

We observed pairs of participants working collaboratively on an analysis task using CoSpaces. Our goal was to gain a better understanding of how people use VRK in a co-located collaborative context and how VRK influences the collaborative analysis process. Therefore, we focus primarily on users' actions that involved history items and notes.

4.1. Participants

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

4.2. Apparatus

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

4.3. Tasks and Procedure

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

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

Transitions between tightly and loosely coupled work happen naturally in collaboration [12],[18]. In practice, participants might know the main analysis task in advance and perform some pre-analysis before the meeting. Later, during the meeting, sharing of their findings and insights would provoke further discussions and raise new questions. At this point, if the problem requires, the team may break up to investigate the newly posed problem, and join together again after finishing the ad hoc analysis. Task 2 tries to replicate as much of this process as possible, but leaves out the pre-meeting analysis for logistical reasons. The motivation behind Task 2 was to ensure we would

observe use of record-keeping under both loosely and tightly coupled collaboration. Data used for this study were sales revenue, margin and quantity sold of clothing items in eight US states for three consecutive years, and consisted of 9 columns and 3273 rows.

4.4. Data Capture and Analysis

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

5. Findings

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

5.1. Actions on History

We identified nine primary actions (Table 1). These can be conceptually placed within the action-tier of Gotz et al.’s [14] categorization of VA activities. Similar to their action-layer, our actions are domain independent and generic and carry some meaningful information regarding users’ intentions (Figure 3).

Some actions such as review note and review history were done to achieve more than one analytic goal. To infer primary analytic intentions, we relied on our observations and action sequences. For actions that appeared to happen for more than one reason, we analyzed pairs of action units. For instance, we observed that in most cases, a Review History action was followed by Reload Chart or another Review History action. Therefore, we concluded that primary user intentions for Review History were to look for a specific chart and to review a number of charts.

Table 1. Primary actions on visual record-keeping and the frequency of each.

Action	Description	#
Reload a Chart	Reload a previously created chart from the history, either the local history or a collaborator's history.	155
Review History	Review charts within the history, either the local history or a collaborator's history.	128
Manual Save	Manually save a chart into the history pane.	102
Delete	Delete a chart from the history pane.	99
Note Taking	Write down notes in the note pane.	77
Review Notes	Review note(s), add to a note, and / or reload a chart linked to a note.	30
View Current Chart	View the collaborator's current visualization through tabs.	11
Copy Local	Copy a chart from a collaborator's worksheet to the local Worksheet.	4
Create External Worksheet	Creating a new Worksheet using a chart from the history.	3

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

Figure 3: Actions on history and their primary user intentions. Right column presents the frequency of each combination.

We recognize that our inferences may not always be correct, and so these numbers should be taken as

approximate. For example, some instances of the Reload Chart action could have been to replace a wrongly reloaded chart. In addition, the frequencies of actions and primary user intentions likely depend somewhat on the system design and the individuals. For example, we suspect that there would have been fewer Delete Chart actions if the tool had used a smarter automatic saving algorithm that saved fewer charts. Nonetheless, we believe that the frequency of such misclassifications is very small, so that the high level trends are still valid.

5.2. Actions and Analysis Phases

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

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

On average, groups spent 75% of the time on information foraging and 25% on discussion. Therefore the number of actions performed during information foraging (455) is almost triple that during discussion (154). Nonetheless, a closer examination of actions performed in each phase shows that there is a relationship between analysis phases and the actions performed on history. As shown in Figure 4, copy local, manual save, note taking and delete actions almost entirely happened during information foraging. Note reuse and view current chart were mostly performed in discussion. Reload a chart, review history and create external worksheet actions were common in both phases.

Of the actions that occurred frequently in both phases, Reload Chart and Review History were the most interesting because users' intentions for undertaking them changed depending on the phase (Figure 5). For example, during discussion, reloading a chart was much more likely to be for the purpose of reviewing it than reusing it for new analysis, whereas

during information foraging, both purposes were common. Also, reviewing the set of previously created charts was much more common during information foraging than discussion, whereas looking for a specific chart was more common during discussion.

5.3. Actions and Collaboration Styles

The design of Task 2 required an “independent, parallel work” style, as defined by Isenberg et al. [10]. In this collaboration strategy, each participant works on the problem independently and results are then combined and discussed at the end.

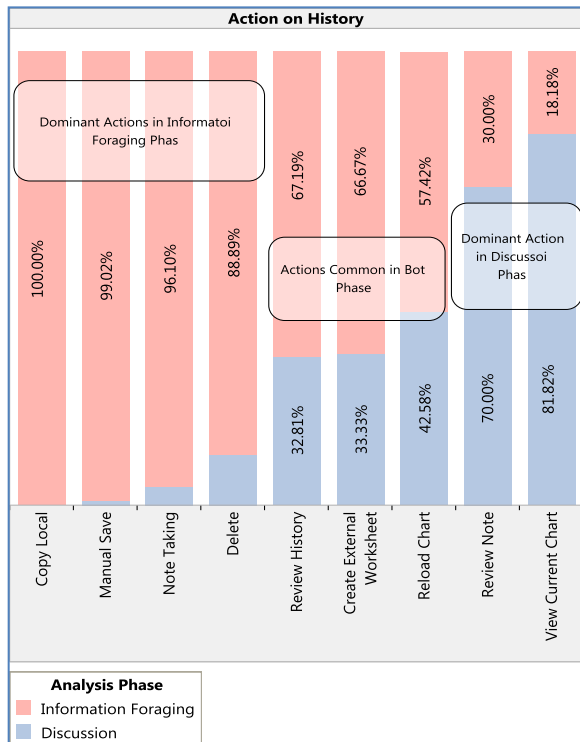


Figure 4: Record-keeping actions in each phase.

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

Figure 5: User intentions for reload chart and review history actions in different analysis phases. Numbers show count of observed actions.

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

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

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

5.4. Record-Keeping Behaviours

Participants used visual record-keeping to capture what they found important in the data, and their explanations for those findings. Based on our observations, we identified two dominant record-keeping strategies. The first was a “note taking oriented strategy”, in which participants took considerably more notes. The second was “chart saving oriented”, in which participants manually saved a larger number of charts. As shown in Figure 6, with the exception of groups 1 and 7, the rest of the groups exhibited primarily one of these behaviours. It is also evident from the same figure that capturing important findings predominantly took place during the

information foraging phase. This makes sense since during this phase participants were investigating data for findings. In the discussion phase only a few notes were taken to record group discussion outcomes.




Figure 6: Count of Manual Save and Note Taking actions by each group. Groups are sorted by strategy.

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




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

5.5. Use of Tabs

Tabs were used to view contents of another worksheet and were used almost equally in both information foraging (15) and discussion phases (17). During information foraging, tab use mostly took place in the middle of the phase, whereas tab use during discussion was almost evenly distributed throughout the phase. This observation matches the logical flow

and objective of the phases. Participants started the information foraging phase by analyzing data and gathering findings and there was not enough history built to motivate remote viewing of a collaborator’s work. As the work in this phase progressed and some history was built up, participants started examining their collaborators’ work. During discussion, participants wanted to share findings and insights; therefore use of tabs to access past items happened throughout the phase.

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

5.6. Quick Review

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




Figure 8: Actions on history while using tabs to view work of the other participant.

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

been performed for making non-detailed comparisons between the current visualization and charts in the history pane.

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

6. Discussion

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

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

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

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

7. Design Implications

A prerequisite for carrying out actions on VRK is proper tool support. There have been general design guidelines for co-located collaborative visual data analysis tools [8],[10],[12],[19], but no specific guidelines for VRK in co-located collaborative VA. Here we suggest some design considerations to improve VRK functionality in co-located collaborative VA on large interactive surfaces.

7.1. Multiple History Views

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

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

Participants in our study copied items from their collaborator's work to use in their own analysis. To facilitate this process, VRK could automatically identify collaborator's records that are related to a user's current line of inquiry. This recommendation could also increase awareness.

In addition to automatic representation of items in history, users should be able to optionally select different history views with varying levels of

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

7.2. Support for Sharing

Our results suggest that changes in collaboration style should be considered in record-keeping design. Participants used tabs under both loose (15) and tight (17) collaboration styles (Section 5.5). Based on this observation, we define "direct" and "indirect" sharing to distinguish between explicit and implicit sharing of work history. In the discussion phase, participants directly shared work to support decision-making. While working individually, participants used tabs to indirectly access their collaborators' work. Therefore, design of an indirect sharing channel should provide remote, unobtrusive and non-interruptive access to a collaborator's work.

One limitation of CoSpaces was that there were no privacy controls. Although this was not a problem in our laboratory study, we anticipate that it could be a concern in real world situations. For example, business analysts in a competitive scenario may wish to carefully guard what they share. Therefore, control over privacy levels [21] and degree of sharing may be a good idea. Users should be able to change the visibility scope of an item among members of the group.

7.3. Support for History Management

Some participants frequently performed manual save and delete actions on the history. At least one participant reported that having to manually delete unwanted items was much more cumbersome than

saving desired ones. This behaviour indicates the perceived importance of history as a container of important analytical artefacts. The observation of abundant "quick reviews" throughout analysis also strengthens this speculation. These results suggest that manually saving items to the history may be a better choice in many circumstances over automatically saving them. Alternatively, an improved algorithm could be devised that would save the analysis state less often and better infer the states that are important to keep. A customizable view (as mentioned in 6.2.1) could also facilitate history management.

7.4. Support for Note Taking & Reuse

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

Participants reported that the link between a note and its related chart was very useful. To further facilitate note reuse, a record-keeping module could automatically create links between related notes based on their content. This would create a network of recorded insights and findings instead of individual notes, which would help analysts to "connect the dots", as suggested by [22].

8. Conclusion and Future Work

We observed nine primary actions on history and key user intentions for each action. These actions and intentions varied depending on the analysis phase and collaboration style. During information foraging, when group work was loosely coupled, history and notes were primarily used to record findings, support individual analysis, and maintain awareness of others' activities. During discussion, when group work was closely coupled, record-keeping tools were primarily used to help present past findings to collaborators and to record discussion results. Interestingly, a remote view of another workspace was useful in both situations, but for different purposes. Based on our findings, we suggest providing various views of recorded material, showing manually saved rather than automatically saved items by default, enabling people to review collaborators' work in an unobtrusive way,

and developing automatic algorithms to better identify which items are related and important to a user's current analysis task.

In future work, we suggest repeating this study for larger groups and different domains. We also plan to further investigate automatic recommendation of history items as suggested by Gotz et al. [20] and linking notes based on their semantic relationships.

9. Acknowledgments

We thank our colleagues in VisID Lab and SAP for their feedback. This research is funded by NSERC and SAP.

10. References

- [1] G. Mark, "The Effects of Collaboration and System Transparency on CIVE Usage: An Empirical Study and Model," *Presence: Teleoperators and Virtual Environments*, vol. 14, no. 1, pp. 1-52, 2005.
- [2] J. Heer, J. Mackinlay, C. Stolte, and M. Agrawala, "Graphical histories for visualization: supporting analysis, communication, and evaluation.," *IEEE Trans. Visualization and Computer Graphics*, vol. 14, no. 6, pp. 1189-96, 2008.
- [3] J. Heer, F. B. Viégas, and M. Wattenberg, "Voyagers and voyeurs: Supporting asynchronous collaborative visualization," *Communications of the ACM*, vol. 52, no. 1, pp. 87-97, 2009.
- [4] F. B. Viegas, M. Wattenberg, F. van Ham, J. Kriss, and M. McKeon, "ManyEyes: a site for visualization at internet scale.," *IEEE Trans. Visualization and Computer Graphics*, vol. 13, no. 6, pp. 1121-8, 2007.
- [5] P. Isenberg and D. Fisher, "Collaborative Brushing and Linking for Co-located Visual Analytics of Document Collections," *Computer Graphics Forum*, vol. 28, no. 3, pp. 1031-1038, Jun. 2009.
- [6] A. Sarvghad, N. Mahyar, and M. Tory, "History Tools for Collaborative Visualization," *Collaborative Visualization on Interactive Surfaces 2009*, p. 21-24, 2009.
- [7] P. Isenberg and S. Carpendale, "Interactive tree comparison for co-located collaborative information visualization.," *IEEE Trans. Visualization and Computer Graphics*, vol. 13, no. 6, pp. 1232-9, 2007.
- [8] J. Heer and M. Agrawala, "Design considerations for collaborative visual analytics," *Information Visualization*, vol. 7, no. 1, pp. 49-62, Feb. 2008.
- [9] J. Heer, F. van Ham, S. Carpendale, C. Weaver, and P. Isenberg, "Creation and collaboration: Engaging new audiences for information visualization," *Information Visualization*, pp. 92-133, 2008.
- [10] P. Isenberg, D. Fisher, M. R. Morris, K. Inkpen, and M. Czerwinski, "An exploratory study of co-located collaborative visual analytics around a tabletop display," *Visual Analytics Science and Technology 2010*, pp. 179-186.
- [11] L. Bavoil et al., "VisTrails: Enabling Interactive Multiple-View Visualizations," *VIS 05. IEEE Visualization*, 2005, pp. 135-142.
- [12] N. Mahyar, A. Sarvghad, and M. Tory, "A closer look at note taking in the co-located collaborative visual analytics process," *Visual Analytics Science and Technology*, 2010, pp. 171-178.
- [13] S. Hunter, P. Maes, and S. Scott, "MemTable: an integrated system for capture and recall of shared histories in group workspaces," *Proc. CHI 2011*, pp. 3305-3314, 2011.
- [14] D. Gotz and M. X. Zhou, "Characterizing users' visual analytic activity for insight provenance," *Information Visualization*, vol. 8, no. 1, pp. 42-55, 2009.
- [15] S. D. Scott, K. D. Grant, and R. L. Mandryk, "System guidelines for co-located, collaborative work on a tabletop display," in *Proc. European Conference on Computer Supported Cooperative Work*, 2003, pp. 159-178.
- [16] Y. Chen, S. Barlowe, and J. Yang, "Click2 Annotate: Automated Insight Externalization with rich semantics," *Visual Analytics Science and Technology*, 2010, pp. 155-162.
- [17] H. R. Lipford, F. Stukes, W. Dou, M. E. Hawkins, and R. Chang, "Helping users recall their reasoning process," in *Visual Analytics Science and Technology*, 2010, pp. 187-194.
- [18] A. Tang, M. Tory, B. Po, P. Neumann, and S. Carpendale, "Collaborative coupling over tabletop displays," *Proc. CHI 2006*, p. 1181-1190, 2006.
- [19] C. Gutwin and S. Greenberg, "Design for individuals, design for groups: tradeoffs between power and workspace awareness," *Proc. CSCW 1998*, pp. 207-216, 1998.
- [20] D. Gotz and Z. Wen, "Behavior-driven visualization recommendation," *Proc. Intelligent User Interfaces 2009*, pp. 315-324, 2009.
- [21] N. Mahyar, a. Sarvghad, and M. Tory, "Note-taking in co-located collaborative visual analytics: Analysis of an observational study," *Information Visualization*, published online Jan, 2012, pp.190-204.
- [22] Y. B. Shrinivasan, D. Gotzy, and J. Lu, "Connecting the dots in visual analysis," *Proc. Visual Analytics Science and Technology 2009*, pp. 123-130, 2009.