

Effects of Implicit Sharing in Collaborative Analysis

Nitesh Goyal¹, Gilly Leshed¹, Dan Cosley¹, Susan R. Fussell^{1,2}
¹Department of Information Science, ²Department of Communication
Cornell University, Ithaca, NY
ngoyal@cs.cornell.edu, {gl87, drc44, sfussell} @cornell.edu

ABSTRACT

When crime analysts collaborate to solve crime cases, they need to share insights in order to connect the clues, identify a pattern, and attribute the crime to the right culprit. We designed a collaborative analysis tool to explore the value of implicitly sharing insights and notes, without requiring analysts to explicitly push information or request it from each other. In an experiment, pairs of remote individuals played the role of crime analysts solving a set of serial killer crimes with both partners having some, but not all, relevant clues. When implicit sharing of notes was available, participants remembered more clues related to detecting the serial killer, and they perceived the tool as more useful compared to when implicit sharing was not available.

Author Keywords

Implicit sharing; collaborative analysis; sensemaking.

ACM Classification Keywords

H.5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

INTRODUCTION

In 2007, Robert Pickton was convicted for six murders of women in British Columbia and connected to 24 others in the Vancouver region [31]. The Vancouver Police Department (VPD) came to suspect that the cases involved a serial killer, but they did not communicate this hypothesis to their cooperating partner, the Royal Canadian Mounted Police (RCMP). Furthermore, missing women reports filed with one agency were not shared with the other except when specific requests were made. These problems potentially delayed the investigation and led to more victims.

In complex crime investigations that span geographical regions, time periods, and cases, relevant information and insights are distributed among individuals, teams, and agencies. These entities need to exchange information in order to create a full picture of the case [14, 27], but barriers ranging from privacy settings, institutional policies to

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

CHI 2014, April 26 - May 01 2014, Toronto, ON, Canada
Copyright is held by the owner/author(s). Publication rights licensed to ACM.
ACM 978-1-4503-2473-1/14/04...\$15.00.
<http://dx.doi.org/10.1145/2556288.2557229>

the costs of exchanging information hinder sharing even when incentives [5] or tools that recommend previously shared information [28] are available.

In this paper we explore whether *implicit sharing*, in which the system automatically shares insights between analysts, can support distributed collaborative analysis. The design idea is that making sharing automatic reduces friction for knowledge sharing, and improves awareness (of otherwise non-shared knowledge) that might limit collaboration, while sharing notes rather than raw data both reduces cognitive load and supports organizational policies around information ownership. We examine the effects of such implicit sharing compared to a system where analysts must explicitly tell their partner about their insights through chat. In an experiment, pairs of remote individuals played the role of crime analysts solving a set of serial killer crimes with both partners having some, but not all, relevant clues. Participants with implicit sharing detected more relevant clues and rated the collaborative features of the tool as more useful, without increasing cognitive workload or reducing explicit communication.

This paper contributes to a growing body of literature in the area of collaborative analysis and sensemaking, by providing findings from a controlled experiment that focuses on the value of implicit knowledge sharing for task performance, use of interface features, team experience, and cognitive workload. Most previous work evaluated collaborative analysis tools with a handful of participants without comparing tools to alternatives [3, 9, 34, 40], by running a solo study [23], by studying explicit knowledge sharing [11, 13], or have not evaluated the tools with human participants at all [30, 32, 35]. In this paper we present a methodical study that demonstrated the benefit of a specific design feature, implicit knowledge sharing, on the process and outcomes of a collaborative analysis task. Understanding the effects of a specific design decision on aspects of collaborative analysis can help researchers and designers make informed decisions before adding specific design features to complex analysis tools.

BACKGROUND: INFORMATION SHARING IN COLLABORATIVE ANALYSIS

In order to analyze and solve crimes, investigators systematically examine timely and pertinent information in search for patterns and trend correlations in the cases they are investigating [16]. According to Pirolli and Card [33], analysts first forage for relevant information in large amounts of dynamically changing data, often from many different

sources. They then schematize the information, creating a mental model that explains the foraged information in a coherent representation. While some of the foraged information might satisfy this representation, other pieces might not fit. Analysts then iteratively forage for additional information to support their mental model or to develop competing hypotheses [20]. Finally, analysts choose the best explanation and schema that fits the information and disseminate it to the relevant audience.

When multiple analysts work together to solve a crime or a set of crimes, the analysis process (foraging, schematization, and decision making) must be performed jointly and strategically to be successful. Analysts working together may be able to collect more data, provide different perspectives on the data, and stimulate each other's thinking processes through sharing knowledge and insights. Such sharing can improve both process and performance in collaborative sensemaking and analysis tasks, measured by accuracy of task outcomes [19, 39] and recall of decisions [13]. Sharing can also help analysts flexibly organize data to discover insights and form schemas [39, 40] or concept maps [15], identify and link evidence from different sources [4, 6], identify entities in the data [9], and encounter otherwise hidden and overlooked connections between pieces of information [1, 9, 11, 25], leading to better decisions [19]. The success of shared workspaces in team performance has been attributed to promoting exchange of information and data with others [19], as a result improved common ground [39] and awareness of the status of the analysis task and others' activities in the task [13, 34].

In order for these benefits to accrue, however, the sharing must actually happen [22, 25]. In general, research on collaborative analysis casts sharing as an explicit process in which analysts consciously choose to share data objects, annotations, and sketches [8, 29]. For instance, Convertino et al. provide each analyst with two workspaces: a private, role-specific workspace and a public workspace available to all analysts [13]. Analysts choose which information pieces and notes to push to the public workspace. Designs that put the burden of sharing on analysts, however, risk under-sharing because of insufficient knowledge of when to share what, failing to understand the importance of sharing, inability to weave sharing into the task at hand, or costs and discomfort in using tools designed for sharing [5, 21].

One way to promote explicit sharing is to motivate sharing through incentives [5]. For example, CrimeFighter [30] supports sensemaking and dissemination by inducing participants to voluntarily share information. Another approach is to remind analysts to share relevant pieces of information with others, as in AnalyticStream [28]. However, even when systems encourage voluntary sharing, organizational norms around information ownership, personal beliefs about these norms, or lack of explicit protocols around information sharing may also inhibit actual sharing [10, 21, 24], as in the Pickton case [31].

Instead of explicit sharing information, we examine the potential value of *implicit sharing of insights*. We focus on insights (e.g., hypotheses about who committed a crime, which analysts often find valuable for joint analysis [3]) rather than raw data and facts (e.g., clues from missing women reports) because the latter is often subject to organizational policies and norms of sharing [10]. Thus, unlike other studies (e.g., [9, 19]) we assume that analysts do not—and should not—necessarily have access to others' raw data sources: separation of information represents organizational boundaries and individual expertise [13]. Instead, as analysts make their own notes about insights from their own evidence, those notes are automatically shared with their partners. This removes the effort involved in assessing whether or not to share an insight and then explicitly placing it in a public workspace [13, 19]. Our hypothesis is that because of this reduction in effort and increase in sharing and awareness, analysts working collaboratively will perform better when implicit sharing is available than when it is not:

H1. Participants using implicit sharing of notes will perform better on a collaborative analysis task than participants without implicit sharing of notes.

Implicit sharing may also shift the value of certain elements of the analysis workspace. Individual features of analysis tools emphasize different aspects of sensemaking, and small changes to these features (such as making them shared vs. individual) can affect analysts' sensemaking strategies [23]. A tool used to capture notes and insights may therefore become more valuable when it is shared with other analysts. Similarly, a workspace viewed only by a single analyst may be less valuable than one that multiple analysts can view. If analysts perceive that these tools are more valuable, they might use them more, interacting with their features and manipulating the data in them. We hypothesize that the availability of implicit sharing of notes will therefore affect both people's evaluations of the features of the tool and their use of these features.

H2a. Participants using implicit sharing of notes will rate the usefulness of collaborative features of the tool higher than participants without implicit sharing of notes.

H2b. Participants using implicit sharing of notes will interact with collaborative features of the tool more than participants without implicit sharing.

We also believe that implicit sharing has the potential to improve the experience of working together. For example, sharing document collections was shown to be valuable to get novice analysts up to speed with the status of what others are doing [3]. In medical settings, implicitly shared awareness information can ease the flow of communication and establishment of common ground between clinical staff members [2, 29]. Similarly, in an emergency management setting, increased common ground was associated with higher perceptions of the team process [12]. If implicit shar-

ing mitigates barriers of collaborating on a task, then individuals should have a better collaboration experience, compared to when implicit sharing is not available:

H3. Participants using implicit sharing of notes will rate their team experience higher than participants without implicit sharing of notes.

By changing the amount and type of information available, implicit sharing may affect the mental demand of the crime-solving task. On the one hand, a shared workspace may reduce the time and effort put in the analysis task compared to working alone [15]. Further, implicit sharing helps establish common ground [39], and thus might reduce analysts' need to explicitly formulate messages and communicate information, thereby reducing their workload. On the other hand, shared workspaces might increase communication costs [19]. Seeing partners' activity might divert attention from one's own thoughts and increase the need for explicit discussion of process and data, especially when shared insights are connected to unshared data [13]. Since the direction of impact is unclear, we pose two research questions:

RQ1. How will implicit sharing of notes affect participants' cognitive workload?

RQ2: How will the availability of implicit sharing affect the amount of information exchanged via explicit channels?

METHOD

We designed a laboratory experiment in which two-person

teams attempted to solve a set of crimes in a simulated geographically distributed environment. The crime cases were distributed between the partners, with a hidden serial killer that had to be identified. Half of the pairs worked on the task using an interface that provided implicit sharing of notes. The other half worked on the task using an interface without implicit sharing. We collected data via post-task surveys, participant reports, and computer logs.

Research Prototype Tool

We adapted Goyal et al.'s SAVANT system [17] for use in the current study. SAVANT has two main components, the Document Space (Figure 1) and the Analysis Space (Figure 2). The Document Space has a number of features for data exploration and discovery. A document library and a reader pane are for viewing and reading crime case reports, witness reports, testimonials, and other documents. A network diagram visualizes connections between documents based on commonly identified entities like persons, locations, and weapon types. The Document Space also provides a map of the area where crimes and events were reported and a timeline to assist in tracking events over time. Users can highlight and create annotations in the text of documents, locations on the map, and events in the timeline.

Such annotations automatically appear in the Analysis Space, an area for analysts to iteratively make and reorganize their notes until they see emerging patterns that lead to hypotheses [23]. These annotations are represented as digital *Stickies* (shaped as a Post-It note, a familiar metaphor

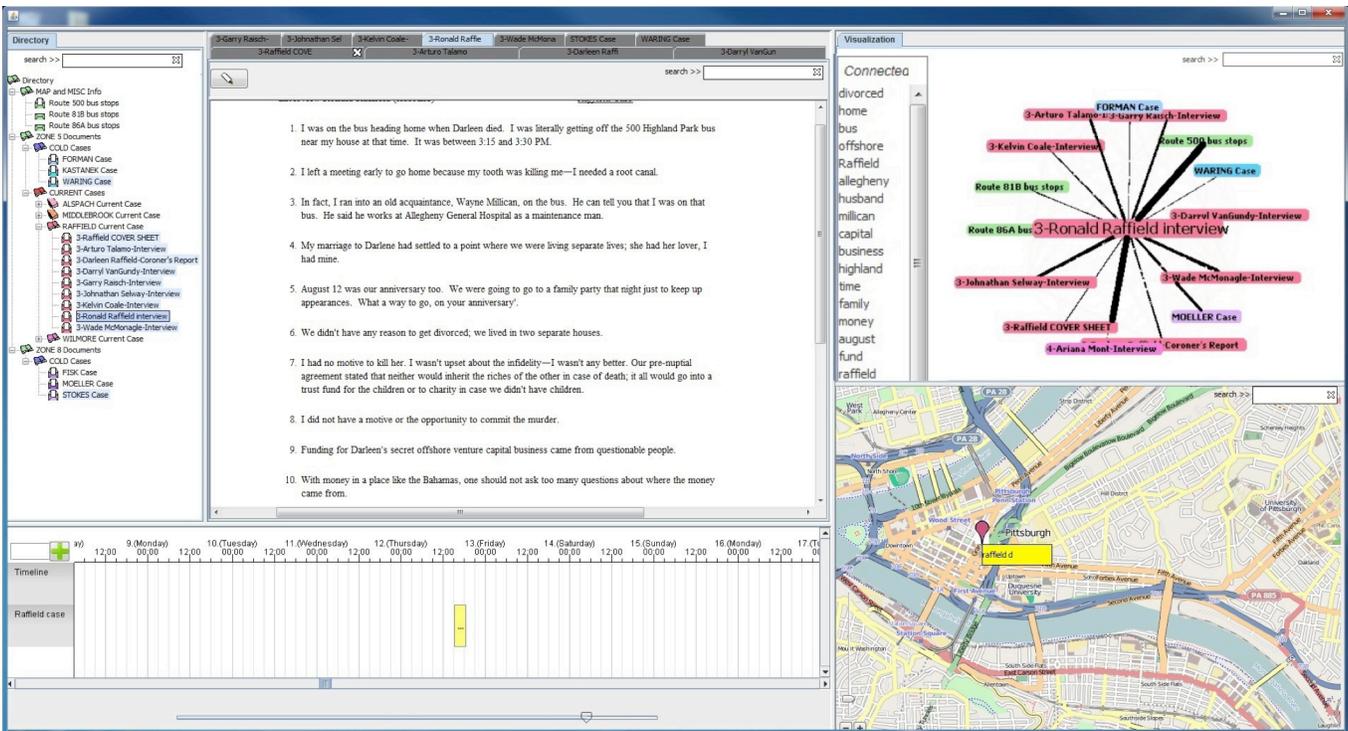


Figure 1. The Document Space showing (clockwise, from top-left) the directory of crime case documents, a tabbed reader pane for reading case documents, a visual graph of connections based on common entities in the dataset, a map to identify locations of crimes and events, and a timeline to track events.

preferred by analysts [7]), which as in other analysis tools [19, 23, 32, 40] are linked to the original document, map location, or timeline point where they were created. Stickies can also be created directly in the Analysis Space, unconnected to specific documents. Analysts can move Stickies around, connect Stickies together, or stack them in piles.

The Analysis Space supports collaboration through both explicit and implicit information sharing. For explicit sharing, a chat box at the bottom-left corner allows analysts to discuss their cases, data, and insights and to ask and answer questions. For implicit sharing, the Analysis Space shows other analysts' Stickies in real time as they are created and organized in the space. Stickies are color coded by the analyst who created them, but anyone can move, connect, or pile anyone's Stickies. Mouse cursors are independent of each other, while dependencies between Stickies are handled by the server on a first-come-first-serve basis. The server updates the interface every second.

We created two versions of SAVANT for this study. In the *implicit sharing* condition, Stickies in the Analysis Space are automatically shared as described above: there is no private workspace for analysis, only a public one. In the *no implicit sharing* condition, partners only see their own Stickies in the Analysis Space: there is no public workspace, only private ones for each analyst. The chat box is available in both conditions to support explicit sharing.

Participants

Participants consisted of 68 students at a large northeastern

US university (22 female, 46 male; 85% U.S. born). Participants were assigned randomly into pairs, and each pair was randomly assigned to either the implicit sharing or the no implicit sharing condition.

Materials

The experimental materials were adapted from Balakrishnan et al. [1] and consisted of a set of practice materials and a primary task. A practice booklet with a set of practice crime case documents introduced participants to the process of crime analysis and highlighted the importance of looking for motive, opportunity, and the lack of alibi. The primary task was created to be reasonable, but difficult, for novice analysts to complete in a limited time. The main task materials were a set of fictional homicide cases. There were six cold (unresolved) cases, and one current (active) case. Each of the cold cases included a single document with a summary of the crime: victim, time, method, and witness interviews. Four of these six cold cases were "serial killer" cases. These four had a similar crime pattern (e.g., killed by a blunt instrument). The active case consisted of nine documents: a cover sheet, coroner's report, and witness and suspect interviews. Additional documents included three bus route timetables and a police department organization chart.

The documents were available through the SAVANT document library and were split between the two participants such that each had access to 3 cold cases (2 serial killer cases, 1 non-serial killer case) and 5 documents from the active case (both participants had access to the cover sheet).

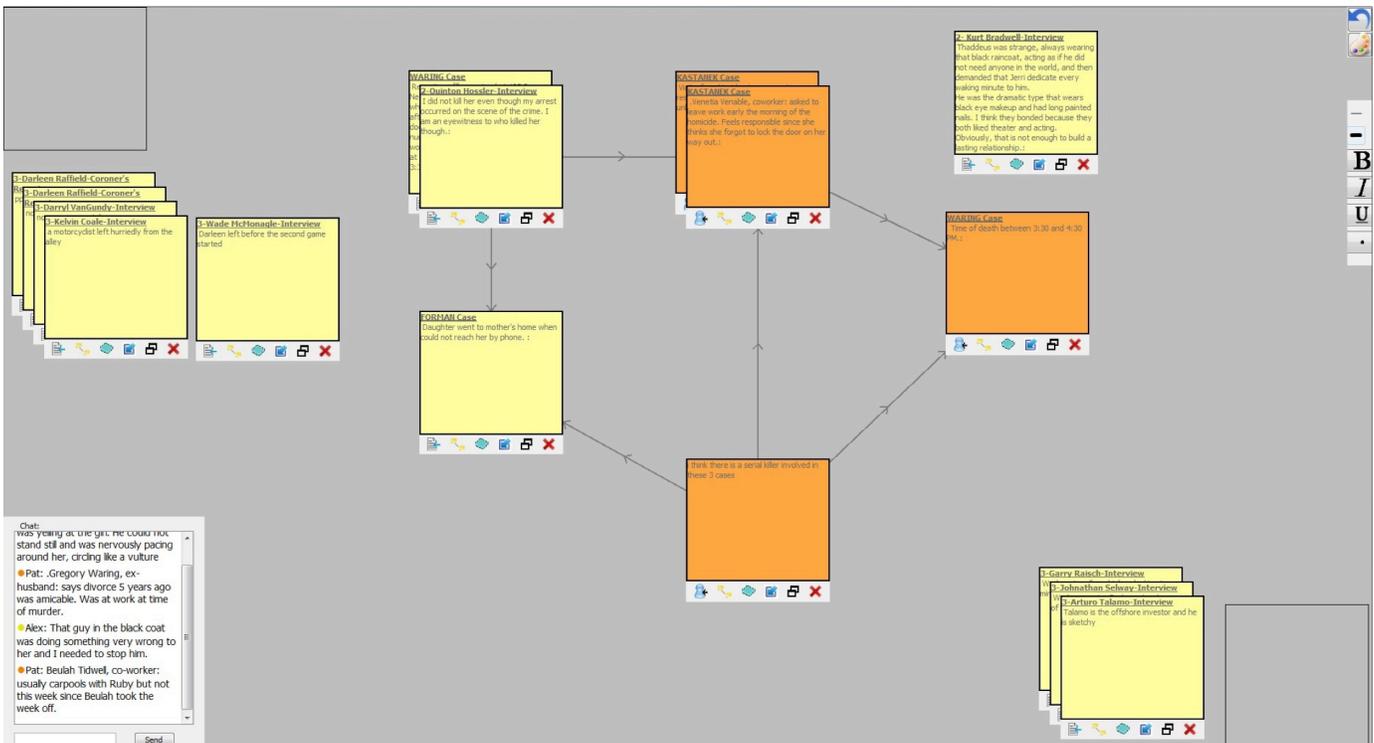


Figure 2. The Analysis Space showing Stickies that are implicitly shared between analysts (color-coded by user), connections between Stickies via arrows, and piles of multiple Stickies. Explicit sharing is supported via the chat box at the bottom left.

The additional documents were available to both participants. Overall, each participant had access to 13 documents, of which 6 were shared with the other participant and 7 were unique.

Twelve clues for detecting the serial killer were dispersed across the 20 documents with 40 suspects/witnesses, equally distributed between the two participants with four in common and four unique to each partner, following a hidden profile task paradigm [36]. The key clue to naming the killer was included in one of the witness reports of the active case, although the active case was not one of the serial killer cases. The task for this study was carefully designed to include data and aspects that are similar to real-world crime cases that remain unsolved, at a scale that could be analyzed in a one-hour session, and at a level of difficulty where many people are unable to solve the crime [1, 17].

Equipment

Two workstations (Intel Core i7 processor, 16 GB RAM) were connected to the Internet and ran SAVANT. Each was connected to two 25" monitors, the left showing the Document Space, and the right showing the Analysis Space. SAVANT logged keyboard and mouse activity as locally stored time-stamped CSV files. To simulate remote collaboration, the workstations were in separate cubicles to prevent eye contact and participants wore noise-cancelling headphones to prevent noises (e.g., keyboard and mouse clicks) from affecting each other.

Procedure

After being seated in separate cubicles, participants signed a written consent form, and read the training materials and performed the practice task individually for about 10 minutes. Participants then received a 10-minute tutorial on the SAVANT interface. The experimenter explained the different parts of SAVANT using example tasks that participants would perform.

Then, using SAVANT, participants worked as a team on the primary task to identify cases associated with a serial killer, name the serial killer, and find as many clues as possible in 60 minutes. At the end of the task, each participant received a paper report form at their workstation to fill out with name of the serial killer, associated cases, and the clues they could recall that would incriminate the killer. They then completed an online survey with questions about clue recognition, the utility of the interface, the collaboration experience, cognitive load, analytic ability (for control) and demographic information.

MEASURES

Our measures are taken from data collected from system logs of interface use, post-task surveys and written reports.

Task Performance

To address H1, predicting that implicit sharing would improve task performance, we used three measures. The first two are based on participants' ability to remember clues

pertaining to identifying the serial killer, and the third is whether they were able to name the killer.

Clue recall. At the end of the session, each participant wrote down as many clues as they could recall supporting their hypothesis about the serial killer. A participant's clue recall score was the number of correct clues written down, similar to the measure used by Convertino et al. [12, 13]

Clue recognition. The post-task survey included multiple-choice questions, each related to one of the 10 clues hidden in the dataset. For example, "On the day of his wife's murder, Ron Raffield claimed that A. He'd run into an old acquaintance on the bus. B. He'd been out of town on a business trip. C. He'd been tied up in a meeting all afternoon. D. He'd tried to call Darlene, but she never answered her cell. E. I do not know." A participant's score was the number of correct answers to these 12 questions.

Solving the case. At the end of the session, each participant wrote a report in which they were asked to name the serial killer. We counted this as binary variable: either the serial killer was identified (1) or not (0).

Perception of Usefulness of SAVANT Features

In order to answer H2a, we asked several questions probing participants' evaluations of features of the SAVANT system in the post-task survey. This is similar to other studies that examined the usefulness of system features [11, 40].

Stickies: Four 5-point questions asked participants about the degree to which the Stickies promoted discussion, helped achieve understanding, and communicate ideas. For example, "The Stickies in Analysis Space helped me understand what my partner was thinking." These four questions formed a reliable scale (Cronbach's $\alpha=0.77$) and were averaged to create a measure of Stickies' usefulness.

Analysis Space. Five 5-point questions asked about the degree to which the Analysis Space helped participants feel physically, cognitively, and emotionally closer to their partner, helped them work with their partner, and helped them understand their partner's activities. These five questions formed a reliable scale (Cronbach's $\alpha=0.85$) and were averaged to measure Analysis Space usefulness.

Use of SAVANT features

In order to answer H2b, we used system logs to derive In order to answer H2b, we used system logs to derive measures of participants' actual use of features in the Analysis Space, including the number of *connections* they made between Stickies, the number of *piles* they created, and the overall number of *movements* (editing, adding, deleting, connecting, or piling) of Stickies.

In the implicit sharing condition participants could manipulate both their and their partner's Stickies, whereas in the non-implicit sharing condition each participant could only manipulate their own Stickies. Therefore, these three measures are at the pair level, aggregating both participants' actions in a session.

Team Experience

The post-task survey contained ten survey questions about the quality of the collaboration (e.g., “It was easy to discuss the cases with my partner,” “My partner and I agreed about how to solve the case”). These ten questions formed a reliable scale (Cronbach’s $\alpha=0.84$) and were averaged to create a team experience score, to answer H3. This measure is similar to [12,13] who used a post-task questionnaire to assess quality of communication within the group.

Cognitive Load

In order to answer RQ1, the post-task survey contained five questions based on the NASA TLX [18] that asked participants to rate how mentally demanding, temporally demanding, effortful, and frustrating the task was, as well as their subjective performance. After inverting the performance question, these five responses formed a reliable scale (Cronbach’s $\alpha=0.76$). Participants’ responses were averaged to create one measure of cognitive load.

Explicit Sharing

SAVANT logged the chat transcripts for each session, which were then cleaned to remove extraneous information like participant identification and timestamps. To answer RQ2, explicit sharing was measured at the pair level as the number of words exchanged in the chat box during a session. This is similar to [19] who assessed the number of chat lines exchanged during the experimental session.

RESULTS

We present our findings in six sections. First, we discuss the effects of implicit sharing on our three task performance measures. We then consider how it affected subjective ratings of SAVANT features, actual use of SAVANT features, perceptions of team experience, cognitive load, and explicit information sharing.

Task Performance

H1 proposed that pairs would perform better when implicit sharing was available than when it was not available. To test this hypothesis, we conducted mixed model ANOVAs, using clue recall and clue recognition as our dependent measures. In these models, participant nested within pair was a random factor and condition (implicit vs. no implicit

sharing) was a fixed factor.

Clue recall. There was a significant effect of implicit vs. no implicit sharing on the number of clues participants recalled in the written report ($F[1, 66]=6.54, p=0.01$). As shown on the left side of Figure 3a, participants in the implicit sharing condition recalled more clues ($M=3.47, SE=0.37$) than those without implicit sharing ($M=2.11, SE=0.37$). Clue recall difference was also significant at the team level ($t[32]=2.03, p=0.05$), with teams with implicit sharing recalling more clues ($M=4.71, SE=0.58$) than teams without implicit sharing ($M=3.11, SE=0.52$). Given the large Cohen’s d (3.68 for individuals, 2.90 for teams) and the fact that these clues were buried in 20 documents with many information pieces, we regard this as a meaningful increase in clue recall, paralleling other work that has found increases in task-relevant recall in shared workspaces [13, 26].

Clue recognition. The right-hand side of Figure 3a shows participants’ performance on the multiple choice clue recognition questions in the post-task survey. There were no statistically significant differences in clue recognition ($F[1, 66]=3.52, p=0.06$) between individuals in the implicit sharing ($M=3.20, SE=0.28$) and no implicit sharing conditions ($M=2.44, SE=0.28$). This was also consistent at the team level ($t[32]=0.80, p=0.42$) between teams with implicit sharing ($M=5.35, SE=0.41$) and no implicit sharing ($M=4.82, SE=0.51$).

Solving the case. We also examined whether interface condition affected the likelihood that participants could solve the crime. Since solving the case was a binary dependent variable, we ran a binomial logistic regression with condition as the independent variable and pair as the random effect variable. There was no significant difference between the implicit sharing condition ($M=0.62, SE=0.11$) and the no implicit sharing condition ($M=0.74, SE=0.01$; Wald Chi Square [1, 68]=0.57, $p=0.45$). Sharing subset of knowledge manually did not improve answer accuracy in [13] but sharing knowledge implicitly in a small experiment did increase answer accuracy in [19].

Perception of Usefulness of SAVANT features

H2a stated that participants would perceive SAVANT fea-

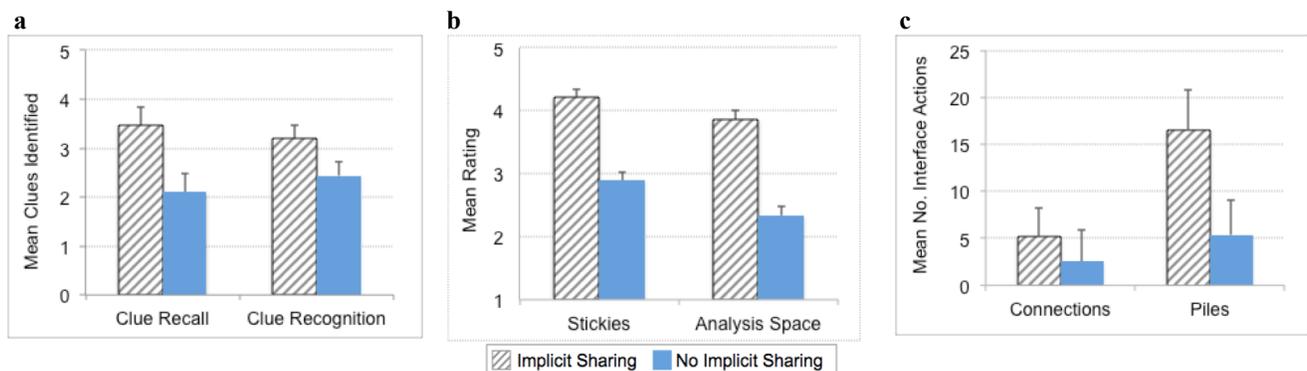


Figure 3. (a) Task performance, (b) Perceived usefulness of Stickies and Analysis Space, and (c) Number of connections and piles made in a session, each by interface condition. Error bars represent standard errors of the mean.

tures as more valuable when the interface supported implicit sharing than when it did not. We analyzed participants' ratings of the usefulness of Stickies and of the Analysis Space using mixed model ANOVAs with participants nested within pair as a random factor and interface condition (implicit sharing vs. no implicit sharing) as a fixed factor.

As shown in Figure 3b, H2a was supported. Participants in the implicit sharing condition viewed Stickies as more useful ($M=4.22$, $SE=0.12$) than those in the no implicit sharing condition ($M=2.90$, $SE=0.12$; $F[1, 66]=53.1$, $p<0.001$). Participants in the implicit sharing condition also rated the Analysis Space as more useful ($M=3.86$, $SE=0.14$) than those in the no implicit sharing condition ($M=2.34$, $SE=0.14$; $F[1, 32]=55.39$, $p<0.001$).

Use of SAVANT features

H2b predicted that the availability of implicit sharing would lead participants to interact more with Stickies in the Analysis Space than without implicit sharing. Using system logs, we counted the number of connections between Stickies, piles of Stickies, and overall Analysis Space manipulations that pairs made over the course of a session. Overall use of connections and piles was quite low and not normally distributed, so we did not perform ANOVAs on this data. Descriptive statistics (Figure 3c) suggest that in the implicit sharing condition, pairs created more connections ($M=5.24$, $SE=3.01$) and more piles ($M=16.59$, $SE=4.16$) than in the no implicit sharing condition (connections: $M=2.51$, $SE=3.38$; piles: $M=5.37$, $SE=3.63$). Participants in the implicit sharing condition also performed more total manipulations of Stickies ($M=1361.85$, $SE=149.02$) than in the no implicit sharing condition ($M=624.79$, $SE=81.95$).

In contrast, Hayne et al. [19] found fewer interface interactions (moving and rating of cards containing information) in a shared workspace compared to an unshared workspace, interpreting this as the need to create and manipulate more individual information elements when there is no access to others' information elements. They conclude that without a shared workspace, the workload on each individual analyst was increased as a result of the increased interactions with interface and information elements.

Team Experience

H3 predicted that participants would rate the quality of their collaborations with their partners higher when they could implicitly share information compared to when they could not. To test this hypothesis, participants' team experience scores were analyzed in a mixed model ANOVA in which participant nested within pair was a random factor and condition (implicit sharing vs. no implicit sharing) was a fixed factor. H3 was not supported ($F[1, 36.30]=0.62$, $p=0.44$).

Cognitive Workload

RQ1 asked whether cognitive workload would vary as a function of the presence or absence of implicit sharing. A mixed model ANOVA showed no significant difference between interface conditions ($F[1, 36.48]=2.49$, $p=0.12$);

participants in the implicit sharing condition rated workload slightly but not significantly lower ($M=4.43$, $SE=0.18$) than in the no implicit sharing condition ($M=4.78$, $SE=0.18$).

Explicit Sharing

RQ2 asked whether the availability of implicit sharing might change the amount of *explicit* sharing via the chat box. A one-way ANOVA was used to compare word counts at the pair level, using condition as the fixed factor. In contrast with [19], who found a significant increase in the amount of explicit chat communication in a shared versus a non-shared condition, we found no significant differences in word count ($F[1, 32]=2.11$, $p=0.16$) between the implicit ($M=782.59$, $SE=77.15$) and no implicit sharing conditions ($M=948.59$, $SE=84.40$).

Roles of Implicit and Explicit Sharing

Participants' open-ended responses on the post-survey shed some light on just how implicit sharing was valuable and how it interacted with explicit sharing features. Several participants mentioned that implicitly shared Stickies helped them "*make connections*" and also added value "*by comparing information*" or "*cross-referencing information*" visually between each other to promote awareness:

"The Stickies enabled a connection between my partner and I, we could see each other's train of thoughts and methods of organization. I used the connecting lines for the Stickies to show myself and my partner the connections that I was seeing." (P27, Female).

Much of the value came from the combination of implicit and explicit sharing. For example, implicit sharing could reduce the need for explicit communication:

"The chat was easily the most helpful because it allowed us to communicate and tell each other specifics about the case. The Stickies were very useful also because they allowed us to make connections between the information we both had independent of talking with each other. [Stickies] allowed us to work more efficiently than wasting both of our time." (P8, Male).

On the other hand, implicit sharing could also prompt explicit chat and sharing, when it revealed needs and gaps:

"I used the Stickies as jumping off points for conversations with my partner - I would see her Sticky and then ask her to fill in some details that she may have skipped over since she had access to certain documents that I did not." (P15, Female).

Finally, Stickies were intentionally designed to be free form and open-ended, and participants did use them to separate aspects of the problem:

"I simply piled them together and placed them in strategic positions. We used two stickies sometimes for the same case. Each sticky would have another side of the case like emotional and the other would be factual." (P61, Female)

“I took notes from the documents and highlighting the important parts created the stickies in the analysis space... I added my notes and thought as stickies in the analysis panes when I made connections between cases.” (P19, Female)

In addition to the quantitative results, these comments show *what* value participants found in implicit sharing, *how* it was used independently and in tandem with explicit forms of communication, and how it can be further improved. They demonstrate the power of implicit sharing to improve collaborative analysis without requiring partners to explicitly push or pull information by triggering understanding and insights on both sides, improving efficiency of conversation, and initiating explicit discussions.

DISCUSSION

Summary of Findings

Our findings show mixed results about the value of implicit sharing for collaborative analysis performance (H1). Participants were better able to identify relevant clues in the data when implicit sharing was available, but were not better able to name the killer. The hidden profile nature of this dataset has been shown to make this task quite difficult [1]. So, although we had hoped to improve outcomes, we are still encouraged by improvements in process elements such as clue recall.

Both participants’ perceptions (H2a) and actual usage logs (H2b) of the Stickies and Analysis Space features of SAVANT suggest that these features were more valuable when implicitly shared—but that they did not increase cognitive workload (RQ1) or change the amount of explicit conversation about the case (RQ2). However, participants’ appreciation of implicitly shared features in SAVANT did not carry over to improved perceptions of the team experience (H3), perhaps because factors such as task difficulty and distributed interaction harmed team dynamics as strongly as the tool helped them.

Limitations and Generalization

One obvious difference between this study and real-world crime investigations is scope. Real investigations have more documents, multiple data types, larger teams, more agencies, and longer durations than our task. This is a general problem for evaluating collaborative analysis tools: the high stakes and established processes of real investigations makes it hard to deploy new tools in the field. Thus, tool evaluations tend to involve short-term tasks completed by university students [1, 11, 12, 13, 19, 23, 39]. These controlled contexts and tasks do provide compensating advantages, such as the ability to systematically vary tools in ways that allow researchers to carefully examine specific factors that together build a body of knowledge to inform real world tools design. Additional research is required to assess how the value of implicit sharing may change over longer time periods, with larger datasets, varied data types, teams, and analysis tasks on factors like cognitive-load etc.

In particular, implicit sharing may not scale to larger tasks if the amount of information to process starts to outweigh the advantages of awareness. Convertino et al.’s explicit moving of information from private to public workspaces [13], Hayne et al.’s idea of parallel personal-but-visible-to-all workspaces [19], and our choice of implicit sharing of all analytical activity are three points in what we see as a large design space for knowledge sharing. Task, scope, team size, and organizational policies might drive the choice of different sharing policies and design decisions (automatic/implicit, manual/explicit, or forbidden/private) for different kinds of information, including raw data, operations on raw data such as reading or organizing it, annotations, and operations on annotations such as connecting and moving them. There are many possible choices around what and how to share that will require a number of rigorous, carefully designed studies that evaluate these specific design features in different contexts.

Finally, the motivation behind our design of sharing insights instead of case documents was to respect organizational policies around sharing confidential materials across institutional boundaries: sharing inferences and notes might allow greater awareness without direct access to confidential raw data. Whether analysts could and would create notes that conform to these organizational constraints, however, is an open question, suggesting the value of careful content analysis (c.f. [13]) in understanding just how the communication features are used and affect sensemaking.

Design Implications

Based on our findings, we suggest a number of design implications that can further improve collaborative analysis performance and process.

One important observation is that not all Stickies were created equal. Our intent was that they would represent “insights” (as opposed to “facts”), but people appropriated the Stickies for a variety of purposes: tracking “emotional” versus “factual” sides of the case (P61); “highlighting important parts” of a case versus “making connections” (P19); “trains of thought” versus “methods of organization” (P27). Providing ways to distinguish between different kinds of analyst note (e.g., through color, font, size, shape) that help analysts bend the notes to their ways of thinking might encourage the sharing of different kinds of information that helps establish common ground [26]. NLP techniques that distinguish document-based facts vs. inferential comments [37, 38] could be used to suggest categorizations, both making this feature smoother to use and encouraging analysts to be more aware of when they are making inferences.

Distinguishing types of notes might also help reduce risk of groupthink that can happen when people are influenced by each other’s behavior. In our case, the overlap between the clues recalled by each partner was high both with implicit sharing (2.23 of 4.71 recalled clues) and without (1.11 of 3.11). This suggests that much of the collaboration focused on already found clues and inferences around them, rather

than finding gaps in the team's knowledge, a robust problem [36]. This might in turn help explain why increased clue recall didn't translate into solved crimes. This is related to the idea of cognitive fixation in brainstorming, where previously expressed ideas in theory inspire new ones but in practice can limit the space where people think.

Thus, providing interfaces that help reduce this tendency become attractive. One option would be to combine implicitly shared, unstructured notes with explicit support for formal hypotheses as provided by CACHE [11], Jigsaw [35], and Sandbox [40]. As with the NLP techniques described above, having semantically different representations for notes and inferences might help analysts separate these activities, reducing the tendency to fixate on hypotheses early and increasing the chance that analysts develop competing hypotheses, which is important to final outcomes [11, 20]. Implicitly shared notes could also be used as resources for creating, supporting, or refuting these hypotheses, with the system suggesting connections between existing notes and newly proposed hypotheses.

Notes also served as a bridge between analysts, according to our qualitative findings about the interplay between implicitly shared Stickies and the chat communication channel (P15). We also propose that they could serve as a bridge between their private knowledge stores: the system could look for similarities between newly posted notes and the documents in each analyst's space. Related documents could be highlighted, helping members of the distributed team do a better job of fact- and hypothesis-checking, sharing of private information relevant to another analyst's work, and doing the kind of cross-referencing that is important in collaborative analysis [23].

CONCLUSION

In this paper, we presented findings from an experiment in which individuals played the role of crime analysts and worked to solve crimes. This study is a step in a series [1, 3, 11, 13, 19, 28, 35, 39, 40] of rigorous, carefully designed studies for testing the impact of specific design features on collaborative analysis processes and performance. A variety of tools have been proposed, designed, and developed for collaborative analysis, but very few evaluate specific design features [1, 19]. Some of these tools have been evaluated with a small number of users and without comparison to alternative tools or to the same tool with design feature variations [3, 9, 34, 40]. Others have evaluated tools in a solo study, failing to examine the collaborative aspects of the analysis task [23], or have not evaluated the tools with human participants at all [30, 32, 35]. Other tools require collaborators to manually decide what and when to share in a structured tabular format [12, 13].

We found that adding implicit sharing of notes to explicit communication via chat improved the ability to detect clues hidden between the partners that represented patterns of a serial killer crime, without sacrificing cognitive workload and explicit forms of communication. We also found no

improved sense of collaboration when the crime analysis system offered implicit sharing of insights with digital notes.

Knowing how implicit sharing may improve task performance is critical to improve the design of collaborative analysis tools between organizations that restrict data sharing for privacy and between collaborators that fail to communicate necessary details. Improving the design of collaborative analysis tools is key to getting more crimes (and other analysis tasks) solved correctly and promptly. We show in this paper that designing implicit sharing into such tools is one step toward these goals.

ACKNOWLEDGMENTS

We thank Wei Xin Yuan and Eric Swidler for their help in developing SAVANT, and Andre Anderson for helping run the study. This work has been funded by NSF grant IIS-0968450.

REFERENCES

1. Balakrishnan, A.D., Fussell, S.R., & Kiesler, S. 2008. Do visualizations improve synchronous remote collaboration? *Proc. CHI '08*, 1227-1236.
2. Bardram, J.E, Hansen, T.R., & Soegaard, M. 2006. AwareMedia: A shared interactive display supporting social, temporal, and spatial awareness in surgery. *Proc. CSCW '06*, 109-118.
3. Bier, E.A., Card, S.K., & Bodnar, J.W. 2010. Principles and tools for collaborative entity-based intelligence analysis. *IEEE Transactions on Visualization and Computer Graphics*, 16(2), 178-191.
4. Boba, R. 2005. *Crime analysis and crime mapping*. Thousand Oaks, CA: Sage Publications.
5. Cabrera, A., & Cabrera, E.F. 2002. Knowledge-sharing dilemmas. *Organization Studies*, 23(5), 687-710.
6. Card, S.K. 2005. The science of analytical reasoning. In Thomas, J.J. and Cook, K.A. (Eds.), *Illuminating the path: The research and development agenda for visual analytics*. *IEEE Computer Society*, 33-68.
7. Chin, G. Jr., Kuchar, O.A., & Wolf, K.E. 2009. Exploring the analytical processes of intelligence analysts. *Proc. CHI '09*, 11-20.
8. Chuah, M.C., & Roth, S.F. 2003. Visualizing common ground. *Proc. InfoViz '03*, 365-372.
9. Chung, H., Yang, S., Massjouni, N., Andrews, C., Kanna, R., & North, C. 2010. Vizcept: Supporting synchronous collaboration for constructing visualizations in intelligence analysis. *Proc. VAST '10*, 107-114.
10. Constant, D., Kiesler, S., & Sproull, L. 1994. What's mine is ours, or is it? A study of attitudes about information sharing. *Info. Systems Research*, 5, 400-421.
11. Convertino, G., Billman, D., Piroli, P., Massar, J.P., & Shrager, J. 2008. The CACHE study: Group effects in computer-supported collaborative analysis. *Proc. CSCW '08*, 353-393.

12. Convertino, G., Mentis, H.M, Rosson, M.B, Carroll, J.M, Slavkovic, A., & Ganoë, C.H. 2008. Articulating common ground in cooperative work: Content and process. *Proc. CHI '08*, 1637-1646.
13. Convertino, G., Mentis, H.M., Slavkovic, A., Rosson, M.B., & Carroll, J.M. 2011. Supporting common ground and awareness in emergency management planning: A design research project. *ACM ToCHI*, 18(4).
14. Egger, S.A. 1998. *The killers among us: An examination of serial murder and its investigation*. Upper Saddle River, NJ: Prentice Hall.
15. Fisher, K., Counts, S., & Kittur, A. 2012. Distributed sensemaking: Improving sensemaking by leveraging the efforts of previous users. *Proc. CHI'12*, 247-256.
16. Gottlieb, S., Arenberg, S.I., & Singh, R. 1994. *Crime analysis: From first report to final arrest*. Montclair, CA: Alpha Publishing.
17. Goyal, N., Leshed, G., & Fussell, S.R. 2013. Effects of visualization and note-taking on sensemaking and analysis. *Proc. CHI '13*, 2721-2724.
18. Hart, S.G., & Staveland, L.E. 1988. Development of a multi-dimensional workload rating scale. In P.A. Hancock & N. Mesh Kati (Eds.), *Human mental workload*, 139-183. Amsterdam: Elsevier.
19. Hayne, S.C., Troup, L.J., & McComb, S.A. 2011. "Where's Farah?": Knowledge silos and information fusion by distributed collaborating teams. *Information Systems Frontiers*, 13(1), 89-100.
20. Heuer, R. J. Jr. 1999. *The psychology of intelligence*. Washington DC: Center for the Study of Intelligence, Government Printing Office.
21. Jarvenpaa, S. L., & Staples, D.S. 2000. The use of collaborative electronic media for information sharing: An exploratory study of determinants. *The Journal of Strategic Information Systems*, 9(2), 129-154.
22. Johnston, R. 2005. *Analytic culture in the U.S. intelligence community: An ethnographic study*. Center for the Study of Intelligence.
23. Kang, Y., Görg, C., & Stasko, J. 2011. How can visual analytics assist investigative analysis? Design implications from an evaluation. *IEEE Transactions on Visualization and Computer Graphics*, 17(5), 570-583.
24. Lee, J., & Rao, H.R., 2007. Exploring the causes and effects of inter-agency information sharing systems adoption in the anti/counter-terrorism and disaster management domains. *Proc. dg.o '07*, 155-163.
25. Mark, G., Carpenter, K., & Kobsa, A., 2003. A model of synchronous collaborative information visualization. *Proc. InfoViz '03*, 373-381.
26. McCarthy, J.C., Miles, V.C., & Monk, A.F. 1991. An experimental study of common ground in text-based communication. *Proc. CHI '91*, 209-215.
27. National Commission on Terrorist Attacks upon the United States, 2004. *The 9/11 Commission Report: Final report of the national commission on terrorist attacks upon the united states*, Norton, NY.
28. Nobarany, S., Haraty, M., & Fisher, B. 2012. Facilitating the reuse process in distributed collaboration: A distributed cognition approach. *Proc. CSCW '12*, 1223-1232.
29. Paul, S.A., & Reddy, M.C. 2010. Understanding together: Sensemaking in collaborative information seeking. *Proc. CSCW '10*, 321-330.
30. Petersen, R.R., & Wiil, U.K. 2011. CrimeFighter investigator: A novel tool for criminal network investigation. *Proc. EISIC '11*, 197-202.
31. Pickton Report:
http://www.ag.gov.bc.ca/public_inquiries/docs/Forsaken-ES.pdf. Pickton Trials:
<http://www.cbc.ca/news/canada/british-columbia/story/2007/01/22/pickton-trial.html>
32. Pioch, N.J., & Everett, J.O. 2006. POLESTAR: Collaborative knowledge management & sensemaking tool for intelligence analysts. *Proc. CIKM '06*, 513-521.
33. Pirolli, P., & Card, S. 2005. The sensemaking process and leverage points for analyst technology as identified through cognitive task analysis. *Proc. ICA Vol. 5*, 2-4.
34. Shrinivasan, Y.B., & Wijk, J.J. 2009. Supporting exploration awareness in information visualization. *IEEE Computer Graphics & Applications*, 29(5), 34-43.
35. Stasko, J., Görg, C., & Liu, Z. 2008. Jigsaw: supporting investigative analysis through interactive visualization. *Information Visualization*, 7, 118-132.
36. Stasser, G., & Titus, W. 1985. Pooling of unshared information in group decision making: Biased information sampling during discussion. *Journal of Personality and Social Psychology*, 48(6), 1467-1478.
37. Stoyanov, S., Cardie, C., and Wiebe, J. 2005. Multi-perspective question answering using the OpQA corpus. *Proc. HLT '05*, 923-930.
38. Wiebe, J., Wilson, T., & Cardie, C. 2005. Annotating expressions of opinions and emotions in language. *Language Resources and Evaluation*, 39(2-3), 165-210.
39. Willett, W., Heer, J., Hellerstein, J., & Agrawala, M. 2011. CommentSpace: Structured support for collaborative visual analysis. *Proc. CHI '11*, 3131-3140.
40. Wright, W., Schroh, D., Proulx, P., Skaburskis, A., & Cort, B. 2006. The sandbox for analysis: Concepts and methods. *Proc. CHI '06*, 801-810.
41. Yi, J.S., Melton, R., Stasko, J., & Jacko, J.A. 2005. Dust & magnet: Multivariate information visualization using a magnet metaphor. *Information Visualization*, 4(4), 239-256.