Cognitive artifacts as a window on design ☆

John M. Carroll *, Marcela Borge, Shin-I Shih

Center for Human–Computer Interaction, College of Information Sciences and Technology, The Pennsylvania State University, University Park, PA 16802, USA

ARTICLE INFO

ABSTRACT

We are investigating *information analysis* as a kind of problem solving in which teams are presented with a collection of facts regarding people, places and events, and then identify underlying connections, patterns, and plans in order to draw specific conclusions. The teams spontaneously created a variety of *artifacts* to hold and organize problem information, and *practices* to simplify and regularize their collaborative interactions around these artifacts. In this paper, we analyze the artifacts and practices as a potential source of insight into how this problem solving activity could be supported by an interactive system design.

© 2013 The Authors. Published by Elsevier Ltd. Open access under CC BY-NC-ND license.

1. Introduction

When people work in complex environments, they support their own intellectual activity by designing *artifacts* to hold and organize problem domain information and *practices* to simplify and regularize interactions. For example, people create descriptive names for their personal files and, in managing file systems, they develop file naming schemas to guide and facilitate the generation of new filenames, and to enhance subsequent retrieval and recognition interactions with files and filenames [5]. This is a simple example of distributed cognition: People creatively shape and leverage the external world to be a more effective resource for their own subsequent activity [22].

Spontaneous ad hoc designs are of course not necessarily optimal designs, or even good designs. Non-designers often inappropriately reuse existing designs [21]. Command languages that users designed to operate robots were used effectively *by those users*, but the user-designed command languages often incorporated linguistic

properties known to evoke command language performance problems for people generally [5].

We are investigating *information analysis* as a kind of problem solving in which teams are presented with a wide array of information regarding people, places and events, and must try to identify underlying connections, patterns, and plans. This is a difficult area to work in because practitioners are often specifically inaccessible to the public because of the security classification of the problem content they address. In a field study of information analysts in the US National Geospatial-Intelligence Agency (NGA, but at that time called the National Imagery and Mapping Agency) only a few participants could be identified, each was only able to devote a couple hours to the study, and all had to be interviewed by proxy, since our team did not have appropriate security clearance to talk to them [27].

We created a *reference task* [35] involving 222 facts pertaining to a set of campus crimes involving stolen laptop computers. We observed teams of three students working to identify suspects, to develop theories of the crimes, and to predict the next likely crime in the series. In addition to problem information, we provided the teams with standard office supplies. Most teams used these physical materials to create ad hoc information artifacts in the course of working on the problem.

In this paper, we describe the artifacts that were spontaneously created, and then consider these ad hoc

---

☆ This paper has been recommended for acceptance by S.-K. Chang.
* Corresponding author. Tel.: +1 814 8632476; fax: +1 814 8656426.
 E-mail address: jmcarroll@psu.edu (J.M. Carroll).

designs as expressions of tool requirements for this pro-blem solving activity. We use this analysis of the artifacts and practices of the student teams to guide a requirements analysis for the design of an interactive system to support the information analysis task. This is analogous to our earlier investigation of emergency management planning [6,28], which we subsequently used to design and evaluate interactive system support for that activity [18].

2. Background

A premise of our work is that studying the artifacts people spontaneously create to support their own activity can be a window into their cognitive and collaborative process. This premise is informed and supported by bodies of research in *distributed cognition*, which regards artifacts people use in carrying out work activity as elements of the overall cognitive system [22,25], in *cultural psychology*, which regards the externalization of thought into the material world as a core strategy for both coping with complexity and for learning [16,33], and in *experience design/embodied interaction*, which regard the materiality of tools and other cognitive resources as critical to their role in human activity [20,36].

Research in cognitive psychology argues that diagrams and graphics are less expressive media than text, parti-cularly with respect to abstractions and relationships, making them easier to understand, often interpretations can be perceived instead of deduced [24,31]. Self-gene-rated external representations function as memory aides, both with respect to the content represented and the person's analysis and interpretation of it [19].

Chin et al. [14] studied five professional information analysts working both individually and as a group. The analysts first made simple annotations in their source materials to highlight purported facts about people, places and events, then they constructed various artifacts to hold and present facts, and finally they tried to identify patterns or connections among facts. The analysts had distinctive and strong beliefs about the artifacts they created and used, and while they believed they could achieve better results by collaborating with others, they also said they would not trust another analyst's artifacts, but would need to review the original source material.

We want to analyze spontaneously created artifacts as a way of gathering and developing implicit design require-ments for technology support. Artifacts that people spon-taneously generate might suggest or embody "natural" ways of representing problem domain information, and therefore might be worth considering as design starting points or design metaphors. More specifically, observed efficacies or difficulties that people experience in using their spontaneous artifacts might suggest approaches to elaborating those artifacts as digital tools [18].

Taking cognitive artifacts as a window on design has known limitations. It is a situated approach, so the artifacts generated in a given problem context may be strongly bound to that context [17]. One way to address this is to try to articulate artifact analyses at a "basic level" of generality [26]. Another limitation is that artifacts spontaneously generated by people may be suboptimal or even poor

representations created to support performance, but per-haps undermining it [15]. One way to address this is to try to link artifact designs with user experiences and perfor-mance outcomes, that is to emphasize representational factors that enhance performance and/or experience, but to mitigate or eliminate those that diminish outcomes for people [7]. For example Schafer et al. [28] showed that maintaining awareness was a key problem for regional emergency planners, which was confirmed and elaborated in laboratory studies. Support for awareness was therefore emphasized in the design of software tools, which indeed did support better performance [18]. Finally, the artifacts people spontaneously generate might function effectively as tools *for those people*, but not necessarily as effective tools for others; this is a version of the "generation effect" [5].

3. Study design

Our task scenario is an analog of the US Navy's Special Operations Reconnaissance (SOR; [34]) scenario in which three information analysts collaboratively synthesize and make sense of a complex information space of people, locations and events. We remapped the scenario content to concern a series of laptop thefts in and around a college campus. This was to better leverage local knowledge of our college student participants (e.g., regarding town and campus geography), and to enhance their engagement in the study.

Each team member was assigned a specialized role with distinct responsibilities regarding information sources: the *Interview Analyst* manages information obtained from interviews with persons of interest (POIs) or witnesses; the *Records Analyst* manages information from reports or files, such as bank/credit transactions, class schedules, police records, etc.; the *Web Analyst* manages information from Facebook, Twitter, EBay, and other online resources. The problem scenario includes 222 pieces of information, or facts, about relevant people, locations, and events regarding the crimes; the 222 facts were evenly distrib-uted among the team members through role-specialized intelligence documents. Participants had to read and analyze these intelligence documents in order to identify the 222 facts.

The mission was organized into three phases, with a specific objective for each phase; the performance scoring rubric is indicated in parentheses. Phase I introduced 105 facts, mainly regarding schedules and relationships of POIs. In this phase, the teams were asked to narrow down a list of 26 POIs to a list of the eight most likely suspects (8 points). *Opportunity* was the main factor for this phase: proximity to the crime scene at the time of the theft.

In Phase II, teams were given 48 further facts. Partici-pants were asked to identify thieves for each of four thefts (4 points), the instigators of each theft (8 points), motives for stealing the laptops (4 points), and whether there were connections among the four thefts (1 point). For this phase, solutions were determined by opportunity and motive: the thieves were one of the eight most likely suspects who either (1) were near the crime scene at the time of the theft and had motive based on their social

Table 1
Schematic overview of the three study phases.

	Phase I	Phase II	Phase III
Information presented	105 facts	48 facts	69 facts
Team objective:	Identify 8 most likely suspects from 26 persons of interest	Identify prime suspects for each of four thefts	Identify suspects for a future crime
Focus of analysis:	Opportunity (time, location)	Means and motives	Opportunity (time, location)
Time constraint:	50 min	45 min	30 min

relationships, or (2) were near the crime scene at the time of the theft and had a relationship with someone else who might had motive. The instigators were the POIs who had motives and were related to someone near the crime scene (means and motives).

In Phase III, teams received 69 further facts about five POIs who are potentially related to a future theft: Participants were told that police had identified a potential victim and four potential perpetrators, based on connections in the analyzed in the previous phases. Participants had to identify schedule overlaps with the victim, and correlate that with access of POIs to relevant map information, to determine which suspect possessed a map with the location of where the victim would be during the time of the potential crime. (For further description of our reference task, see [9]).

Teams worked under time constraints: each team had 50 min for Phase I, 45min for Phase II, and 30min for Phase III. At the end of each task phase, team members were to come to a joint decision and write down a team recommendation or answer for that phase. Table 1 schematically overviews the study procedure.

4. Methods

Participants were 66 undergraduate students recruited from a large northeastern university in the United States. In exchange for their participation, each student earned class credits or received cash. Average age of the participants was 21.3 years; 53% of the students were male. The students represented diverse disciplines, with 28.8% from Information Sciences & Technology, 12.1% from Psychology, and the remaining students (59.1%) from 26 other majors, such as Security and Risk Analysis (4.5%), Supply Chain (3.0%), Mechanical Engineering (3.0%), Communication Arts and Science (3.0%), Advertising (3.0%), and Anthropology (3.0%). Information Science and Technology and Psychology are themselves rather diverse curricula. We treated academic major of the participants as a random variable. Participants were randomly assigned to 22 three-member teams.

Upon arrival, each participant was randomly assigned one of three analyst roles. Participants were told they would play the criminal investigation role of an expert in a specialized information analysis team, and were each given a general Mission Statement and a Phase I Role Document. The Mission Statement explained the crime scenario and the team's overall objectives, as well as background information on criminal investigation [1,13]. For example, teams were told to consider the motives, means, and opportunity of POIs. The Phase I Role Document described that member's team role, including the information sources to be managed by that role in the collaborative activity. The balance of the Role Document provided each of the three participants information pertaining to the problem scenario that was specific to his/her role. Participants were not permitted to exchange their raw information, but were told that to be successful, teams had to share information and make decisions together, that each member of the team held important and unique pieces of information, and that each member was responsible for sharing particular pieces of information when relevant. Phases II and III also were initiated by presenting each team member with their Role Document.

A wide-angle video recorder with tabletop microphone was set-up to capture the entirety of the lab portion of the study, from the experimenter instructions to behaviors of all team members as they completed each phase of the scenario. At the end of each problem phase, teams were given an intelligence update containing the correct analysis and solution for that phase. Thus, each team entered each succeeding phase with the same prior information and analysis.

In order to contrast the characteristics of high and low performing teams, we calculated each team's performance scores for each phase (points obtained/points possible) and selected from the 22 teams the five consistently highest performing teams and the five consistently lowest performing teams for further detailed analysis. In the balance of this paper we describe and analyze data from these ten teams.

Videos of team problem solving interactions for each of the ten teams were transcribed. We employed interaction analysis [23]. Each new speaker utterance and or behavior was parsed as a new turn, and time-stamped. Turns were split into dialogue acts, separate sentences. Nonverbal gestures and events such as moving to the whiteboard, nodding, using hands, and creating artifacts were indicated in the transcript. Any time participants, used, referred to, or wrote on one of their artifacts, this was coded in the transcript. Research group discussions were used to verify observations related to artifact use and organize analysis.

5. Individual artifacts

Participants – individually and collaboratively – created a variety of information artifacts as they studied, organized, analyzed, and draw conclusions from the scenario. We classified the artifacts created into six types: text anno-

tations, calendar annotations, geo-spatial annotations, lists, tables, and graphical representations. *Text annotations* were notes or highlighting inscribed directly on problem documents. For example, many teams wrote on the intelligence updates given to them at the end of each scenario phase with final police reports presenting the correct interim analysis for the crime. *Calendar annotations* were monthly/weekly calendar/planner photocopies that participants wrote annotations on. *Map/geo-spatial annotations* were artifacts that used a provided paper map of the campus as a foundation for further annotations. *Lists* were one-dimensional text-based representations, for example, an enumeration of motives, or of information relevant to a particular crime date. *Tables* were two-dimensional textual representations, a row by column format, in which the majority of information in the cells was textual or numeric (i.e., little to no graphical illustration). For example, a list of suspect rows, crossed by a series of crime date columns, with each cell indicating whether that particular suspect was near a crime scene on that date. *Graphical representations* were composite artifacts, such as calendars with embedded tables, and non-orthogonal illustrations, such as social network graphs.

Some of the information artifacts created by participants were never shared among team members. These individual artifacts were created and used by only a single member of a team. In total, 108 individual artifacts were created across the 30 members of the 10 teams; 29 participants created at least one. The individual information artifacts tended to be relatively simple: There were 46 text annotations (44 made to intelligence documents and two made to mission statements) and 43 lists versus six calendar annotations, two geospatial annotations, nine tables, and two graphical representations (see Table 2). They tended more to be created early in the problem-solving scenario: 55 in Phase I, 27 in Phase II, and 24 in Phase III.

Because the creation and use of these artifacts was inscribed in the activity of a single participant, they left few traces in the video and transcription data. As far as we can tell, these artifacts were created as participants read through problem documents, or as they listened to other team members speak, as personal note taking. For example, 28 of the 43 lists (65%) were names of and information pertaining to persons of interest (POI). These lists were somewhat opportunistic; names appeared in their order of occurrence in the problem documents, and with a wide variety of associated information, including no associated information. In some cases, it appeared that the participant's strategy for creating the artifact had changed part way through its creation; for example, names early in the list included associated information, but names later in the list did not.

Another typical sort of list we identified was one enumerating all of the POIs on one page, but with no associated information. Half of these lists included highlighting or names crossed out, perhaps to help the participant keep track of which suspects had been discussed or eliminated during the course of the team discussion. Fig. 1A illustrates a list artifact of this type. Fig. 1B is an example of a text annotation artifact; in this case, a

Table 2
Distribution of individual and team artifacts into six categories

	Individual artifacts	Team artifacts
Text annotation	46	2
Calendar annotation	6	2
Map/geo-spatial annotation	2	1
List	43	15
Table	9	10
Graphical representation	2	5
Total	108	35

Records Analyst used color-coded highlighting to associate suspicious account transactions with crime dates.

The individual artifacts suggest that team members focused on concrete details versus overall framing information. Twenty-five out of the 30 participants (more than 82%) made annotations to the intelligence documents that contained information about the victims and persons of interest. In contrast, only two out of the 30 participants (less than 7%) made annotations to the mission briefing document that contained information about the problem as a whole, i.e., task goals, verifiability of information, total number of persons of interest, and objectives of the information analysis activity.

6. Team artifacts

The key distinction between individual artifacts and team artifacts is that the latter, but not the former, were used collaboratively, that is, shared among team members. The majority of team artifacts (27/35, or 77.1%) were created as a means to integrate information provided by more than one team member for the team.

All of the team artifacts were created, or at least initiated by a single team member. Unlike individual artifacts, the creation of team artifacts was always preceded by some public behavior, though in most cases this was quite limited, often a single utterance, in which the artifact creator reflected on types of information held by different team members and/or types and relationships in information that might help the team develop its analysis. For example, immediately before creating a relationship diagram (see Fig. 6), a Web Analyst (team 2) stated, "So we need what? Motive, who the thieves are, why they did it, and possible connections?" Typically, this initial interaction regarding team artifacts was essentially just a notification; the artifact creator moved ahead without waiting for a response from the other team members.

The initial planning appeared to function as a public analysis framing which the artifact itself embodied. Team artifacts were also somewhat more refined than individual artifacts. In another example, a team was discussing the concept of opportunity. One member embodied part of the discussion as a person by crime date matrix, with an alphabetical list of POIs as the ordinate, the four crime dates as the abscissa. Each cell of the matrix analyzed the opportunity of a given POI to commit a given crime. This design is quite refined and also quite simple (see Fig. 2C); it was very useful to the team in collating and information
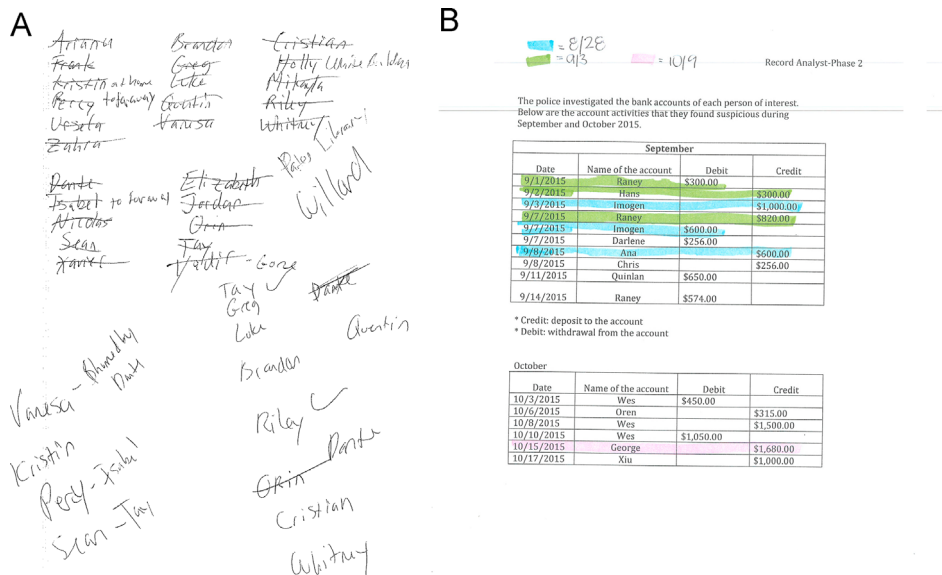
Fig. 1. Individual artifacts (A) (left) is a set of lists enumerating POIs (B) (right) is color-coded text annotations to associate suspicious transactions with crime dates. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)
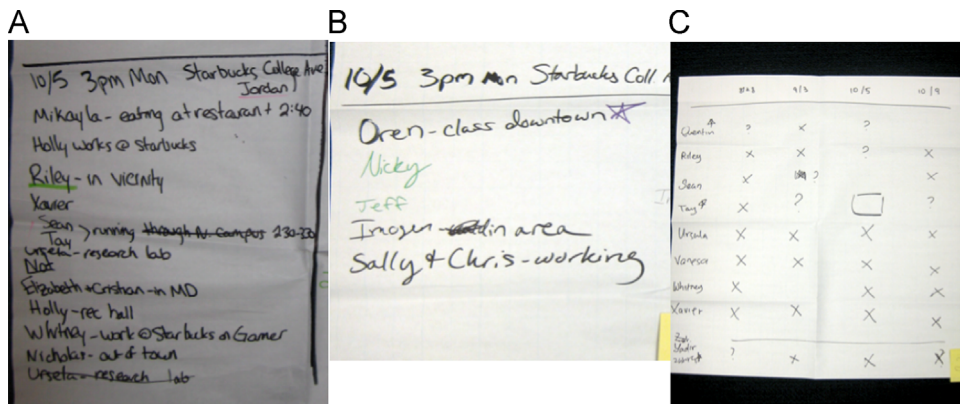

Fig. 2. Artifacts created by three teams in phase I of the task scenario. Example A is a list created by team 17; example B is a list created by team 7; example C is a table created by team 2.

from various problem documents, in organizing the team's discussion and identifying questions they needed to address, and for developing further team artifacts.

In contrast to the individual artifacts, which tended to be created early in the problem scenario (more than half were created in Phase I), team artifacts tended to be created later in the activity: 9 in Phase I, 13 in Phase II, and 13 in Phase III. The maximum number of shared artifacts created by any team was 9 (team 8). Most teams created multiple artifacts; three teams (1, 14, and 15) created a single artifact. Overall, there were two text annotations, both made to intelligence updates, highlighting the correct answers for a previous phase, two calendar annotations, one geospatial annotation, 15 lists, ten tables, and five graphical representations (refer to Table 2).

Team artifacts tended to be more complex: where half of the individual artifacts were direct annotations of problem materials, only 5 of 35 team artifacts were annotations; conversely, where only 10% of the individual

artifacts were tables or graphical representations, for team artifacts this proportion was greater than 40%.

As for the individual artifacts, the team artifacts focused on organizing and emphasizing problem content versus overall problem framing information. Only one team annotated (highlighted) a critical piece of information in the mission briefing regarding the reliability of POI information.

Because the team artifacts involved some public planning (albeit limited) and collaborative interaction as they were used, we were better able to track and analyze when and how they were created and used in video recordings and transcripts. In the next few sections we focus on the team artifacts.

7. Artifact categories

The most typical team artifacts were lists and tables. Teams used lists to represent information related to specific crime dates or motives. Examples A and B in

Fig. 3. A relationship diagram developed by team 18 using post-it notes.

Fig. 2 illustrate the use of lists. Example A is a raw list of all the information the team had about all persons of interest (POIs), including where the POIs were, and relationships among POIs (coded by color). Example B is a filtered list, in this case enumerating only the POIs whose locations gave them an opportunity to commit the crime.

Example C is an example of an elimination table with POIs listed as rows and crime dates as columns. Team 2 inscribed a question mark in a corresponding cell if the given POI had opportunity to commit the crime and an X if he/she did not. This is a highly refined artifact, abstracting the narrative details of specific POIs with regards to their proximity and availability during crime events.

Graphical representations are of interest because they are more complex and often more refined. One team (18) devised a paper artifact that afforded editing (see Fig. 3). This artifact employed post-it notes labeled with the names of POIs, and clustered into affinity groups based on the degree of interpersonal relationships between individuals. This approach provided dynamic flexibility: the post-its could be moved around to reflect further information and analysis. Also, key facts associated with individuals could be written directly onto their corre-sponding post-it. The team used light green post-its to label characteristics of clustered POI post-its. Fig. 3 illus-trates some of the boundary conditions and limitations of this artifact: Some post-its covered others, thus the relat-edness of two post-its is depicted by occluding informa-tion that could be important. Although the POI names are written in large bold lettering, once the post-its have been grouped, it can be difficult to retrieve information by POI name (versus, say, an alphabetical list).

Team 21 devised a grpahical artifact incorporating geo-spatial annotation and two kinds of color-coded post-its on a campus map (Fig. 4). This map was used to depict where people were during each of the crimes, and if they were close enough to commit the crimes. The blue post-its highlight the four crime locations, each labeled in a unique color of ink. The pink post-its indicate the location of various POIs, with the color of ink indicating which of the four crimes they most likely committed. This artifact was created by the team to visually integrate information they
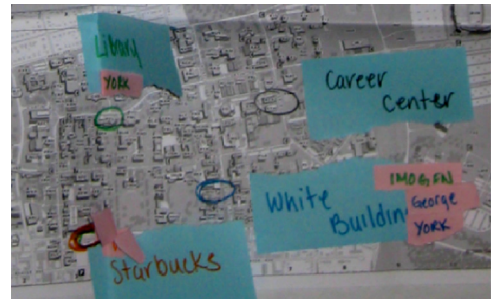

Fig. 4. Team 21 used written annotations and two types of colored post-its on a campus map. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)
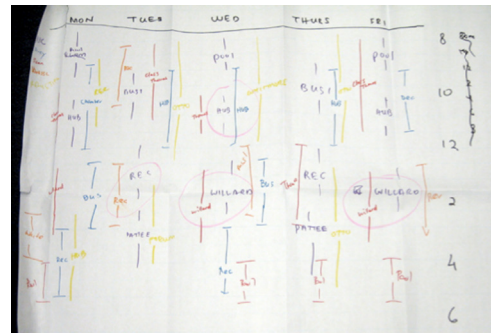

Fig. 5. Diagram developed by team 2 in phase 3 of the scenario to compare schedules across five suspects.

had already collated in a simpler list artifact: suspect names, locations, and crime dates. The team wanted to look at the distance between locations to more directly search for patterns of connection among suspects. One problem they seemed to encounter was extending the color coding of the four crime locations to color coding of the POIs. The main problem they experienced is that they used color coding as a means to highlight too many different things and found it difficult to consistently apply the color coding rules. They followed an abstraction approach to creating the diagram by converting a great deal of textual information into abstract representations (i.e. color codes, stars, "I" for instigator) and then pro-ceeded to add more and more of these these types of represenations to the artifacts.

During Phase III, Team 2 created a composite calendar diagram incorporating the weekly schedules of the five people they were considering—four suspects and the likely next victim (see Fig. 5). This artifact was based on an individual artifact created by the Record Analyst. She suggested developing a comprehensive version incorpor-ating everyone's schedules, and began by placing a large piece of paper on the table and writing the days of the week in black marker accross the top of it. One of her teammates then began to read off his schedule informa-tion, but the third member interrupted to suggest that they use different colored pens to code the schedules of each of the five people. As they worked, the team elabo-rated their approach to the artifact to address various issues. For example, they began by color coding the line segments representing spans of time corresponding to

Table 3
Interaction excerpt: team 18 re-identifying the rationale for including George in a decision table.

Record Analyst: We put George in at Thursday. Why?
Web Analyst: Any reason why?
Interview Analyst: Yeah He has a doctor's appointment, which is right next to [the Career Center …
Records Analyst: We have nothing else that can place him in any other. And thirty minutes before that it happened. The appointment is.]
Web Analyst: Yeah, fine, fine. I get it.

scheduled events for that POI, but realized they needed location information. Accordingly, they inserted abbreviated place name for each schedule event. They also began with the strategy of sharing all schedule information for all POIs, but then realized that they only needed to identify schedule overlaps with the potential victim's schedule. Consequently, they focused on representing the victim's schedule information, and filtering other schedule information depending on whether it overlaped with the victim. Finally, as they started to run short of time, they recognized the inefficiency of having one scribe, and shifted to adding information to the artifact simultaneously. This artifact is very appropriate for Phase III. Although it like other evolving artifacts, it ended up somewhat cluttered, it provided a prompting framework that elicited all the relevant information from the team.

8. Grounding artifacts

All categories of artifacts externalized information from the various intelligence documents presented to members of the teams. They summarized and integrated problem information, and facilitated sharing and pooling information with fellow team members. However, creating artifacts also removed problem information from its original context—the intelligence documents. To some extent, the collaborative team work helped to provide grounding for information in the artifacts. For example, as someone's name, or a short summary of a fact about someone, was inscribed in a list or other artifact, one of the team members might state the rationale for including that person or that fact, as suggested in the descriptions of Figs. 2–5. This is a simple example of fixing reference [5] in which the speech context that accompanies a baptismal use of a linguistic expression is evoked in successive uses of that expression.

Fixing reference is a logical–historical relationship that is often undermined by the limitations of human memory and attention. People can baptise an expression, but then quickly forget what it referred to. Our participants often realized they could not quite recall why a particular name or fact appeared in an artifact, and had to engage in explicit grounding work in which they redeveloped the rationale, as illustrated in Table 3.

The team had already gone through this rationale when they initially included George in the decision table. But as they began to use the table and make their decisions, two of the three members had forgotten the original rationale. This grounding process, or in this case re-grounding, can be seen as a *cost* of artifact abstraction: The rationale for including George did not itself appear in the artifact, indeed were all of the possibly relevant context and

rationale explicitly included the value of the artifact as a summary representation would be compromised by additional visual complexity and clutter.

The re-grounding process can also be seen as a potentially useful process in that it motivates re-examination of the original rationale, which could have been flawed (as well as being forgotten).

9. Artifact strategies

We identified strategies in the creation and use of artifacts by examining videotapes and transcripts of episodes. We drew a distinction between accretion and filtering strategies. *Accretion* is successively adding information into an artifact. The artifact collates and presents problem information more densely, abstracting the information from the original context of the intelligence documents provided in the problem brief. Accretion produces fairly comprehensive artifacts. One of its strengths is that it is relatively mechanical note taking. This makes it easy to coordinate as teamwork. A simple example of accretion is the list in Fig. 2A, enumerating every fact the team had identified about all of the POIs for one of crimes.

Filtering is selective adding, and therefore also excluding of information from an artifact. Filtering abstracts problem information not only from the context of the problem documents, but from the context of other facts judged irrelevant to the purpose of the artifact. Filtering was relatively more interactive among team members; it involves negotiation and reaching consensus, and then monitoring execution. A simple example of filtering is Fig. 2B, enumerating facts involving only the POIs whose locations gave them an opportunity to commit the crime.

Table 4 presents an interaction excerpt in which members of Team 7 recognize the utility of a filtered list of suspects as they create and develop the artifact in Fig. 2B. The Interview Analyst suggests partitioning the list by crime date and location, and making notes person by person. Several minutes later, the Web Analyst suggests that they filter the list, including only likely suspects for the crime. The result is a less rich, but more refined problem representation.

The fundamental tradeoff between accretion and filtering is that accretion is relatively easy to manage and produces authoritative artifacts, in the sense that they are complete representations of the problem data. However, accretion can also produce cluttered displays that are difficult to use. Filtering produces representations that are specially suited to particular issues; artifacts that are simpler and easier to read. Filtering can simplify decision making to a merely perceptual task. Thus, the table in Fig. 2C allows one to see at a glance that several people

Table 4
Interaction excerpt: team 7 recognizing the utility of a filtered list of suspects.

Web Analyst: (Creation of artifact in Fig. 2B) I think it'd be easier if we had a list that we can look at, at each others' stuff, if we had like one list with different suspects and then we can talk about like…
Interview Analyst: Yeah let's just divide that paper by the date of the crime and the place, and write down notes for the person and stuff.
Records Analyst: For each person?
Interview Analyst: Yeah
(Participants continue to share info for 3.23 min until Web Analyst suggests modifying rules for the artifact.)
Interview Analyst: Bob has an internship in DC, and he's been there since August 19th
Web Analyst: So he's not like a suspect for this one. Maybe we should just keep suspects in this stack. You know what I mean?

cannot possibly have committed crimes. Moreover, the teamwork involved in making filtering decisions about what information is relevant, given a filtering strategy, appeared to be engaging to team members. It evoked discussions creating further cognitive retrieval cues for subsequently using the artifact, as illustrated in Table 4. However, constructing a filtered artifact is more difficult, since judgments about information to include or exclude must be made; indeed, it is not just *that* the team can become more engaged, but also that the team *must* be more engaged for filtering to succeed. Finally, filtering raises the possibility of errors that might be hard to recover from, for example, if a relevant piece of information is accidentally excluded in creating a filtered artifact.

Although accretion and filtering are distinguishable strategies, they were rarely used monolithically. Through the course of developing and using artifacts, teams could use both strategies. We classified artifacts as to their primary strategy: recording all shared information versus actively deciding to include/not include shared information. For example, the raw list of Fig. 2A is filtered in the sense that it includes only facts relevant to one of the four crimes. However, the team directed most of their interaction regarding this artifact amassing all information about that crime date, rather than actively evaluating information, based on opportunity or other filtering criteria.

Most artifacts could be classified with respect to primary strategy, but some could not. For example, the graphical representation in Fig. 5, as described above, was initially pursued as a comprehensive representational framework for accretion of all the schedule information presented in Phase III. This artifact is relatively sophisticated, but the original goal was to comprehensively aggregate problem information. However, the team eventually recognized that they did not really need a comprehensive representation, that they could exclude schedule information that did not comprise an overlap with the schedule of the potential victim. At that point, they switched to a filtering strategy, evaluating each piece of information they subsequently included. Accordingly, we classified this artifact as "major strategy not clear."

We analyzed 54% (19/35) of the artifacts as primarily implementing an information accretion strategy, and 31% (11/35) as primarily implementing a filtering strategy. For five of the artifacts (14%), we were unable to categorize the primary strategy as being one or the other ("major strategy not clear").

Most of the artifacts teams created were directly grounded in the problem documents, even if this grounding was actively modulated by a filtering strategy. However a few artifacts appeared to be derived from other artifacts. For example, team 21 created a list identifying all of the suspects for each crime date, distinguishing crime dates with a unique color coding. The team later created a successor artifact from this list: the graphical representation in Fig. 4. This heavily annotated map replicated some of the information from the prior list artifact, and reused the color-coding idea, but specifically helped the team to visualize where POIs were with respect to the crime scenes.

10. Coordinating artifacts

Thus far we have primarily focused on the development and use of separate artifacts, but most of our teams developed multiple shared artifacts. These artifacts often emphasized complementary aspects of the information space, illustrating a higher-order filtering strategy in which the filtering was coordinated across a collection of artifacts. Sets of complementary artifacts often were used together. We analyzed use of 63% (22/35) of the team artifacts as coordinated with other artifacts.

Lists of information were often used with other forms of representations such as tables or diagrams. Team 2 created an exhaustive alphabetical list of all the POIs. They used this artifact primarily to structure their consideration of information with respect to POIs. They referred to the list to help keep team discussion focused on a particular POI, and then successively moved down the list to ensure that they systematically considered everyone. When a teammate moved on prematurely or digressed, the team member in front of the artifact would point to the list or refer to where they were in the list to refocus the discussion.

Fig. 6 depicts the coordinated use of four artifacts (A1, A2, A3, and A4) during a 3-min segment of interaction focused on the second crime (towards the end of Phase II). In the figure, red circles are used to emphasize information in an artifact that the team seems to be discussing. The team considered a suspect, Isabel. Using the table A1, they recognized that they had eliminated Isabel for all the crimes, based on opportunity. The Web Analyst reminds them of the possibility that Isabel facilitated the crime, and they examine the relationship diagram A2, to see who was connected to Isabel ("Isabel" is in the intersection area of the two red circles). They identified Sean, her brother, and Luke, a friend. They then examined the table A3, observing that Sean had a question mark for crime 2, indicating possible opportunity to commit the crime (he was scheduled to be swimming at a location near the crime scene).
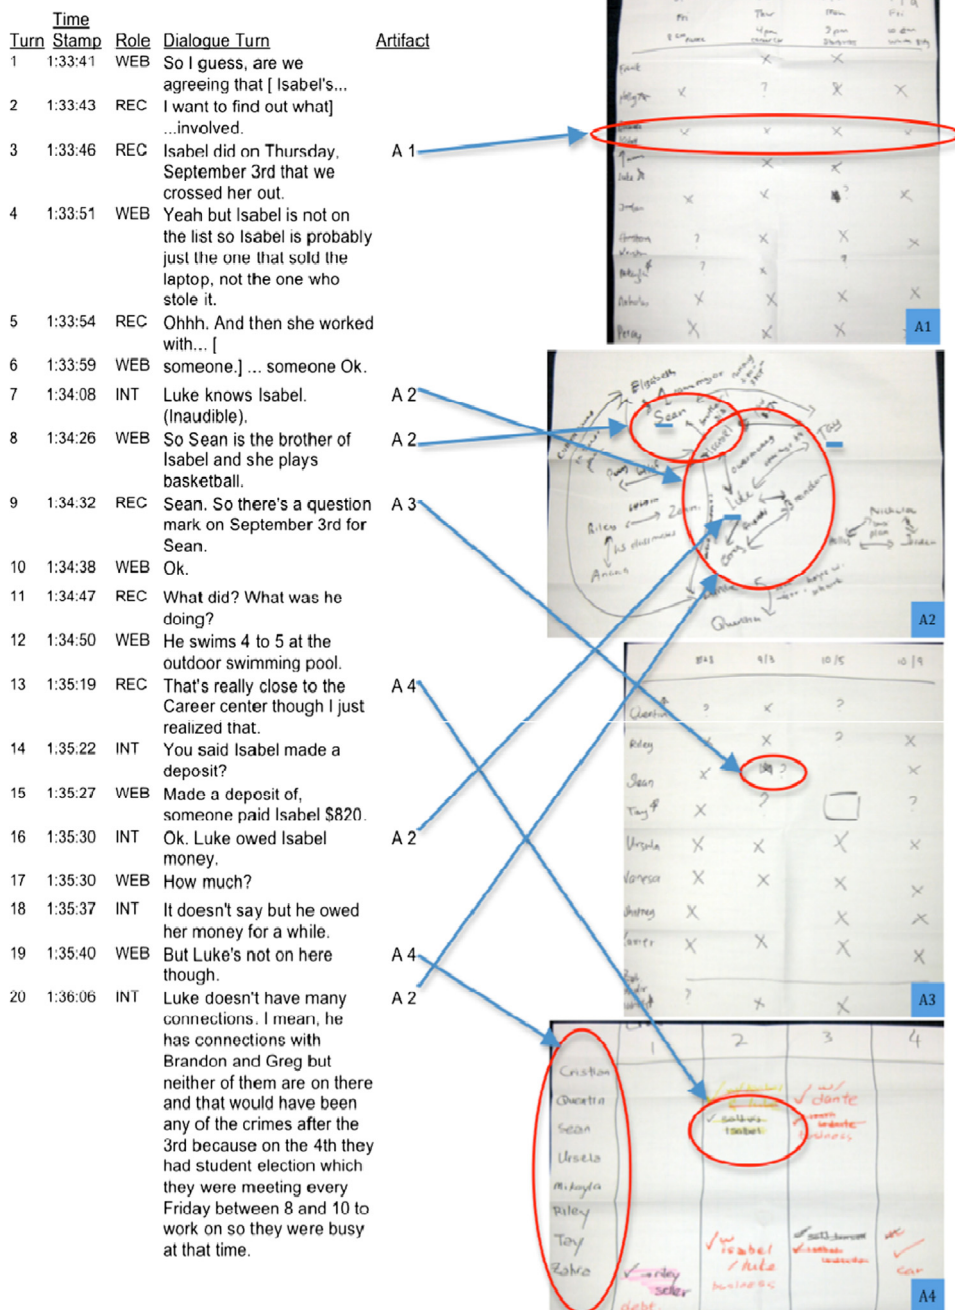
Turn	Time Stamp	Role	Dialogue Turn	Artifact
1	1:33:41	WEB	So I guess, are we agreeing that [Isabel's...	
2	1:33:43	REC	I want to find out what] ...involved.	
3	1:33:46	REC	Isabel did on Thursday, September 3rd that we crossed her out.	A 1
4	1:33:51	WEB	Yeah but Isabel is not on the list so Isabel is probably just the one that sold the laptop, not the one who stole it.	
5	1:33:54	REC	Ohhh. And then she worked with... [	
6	1:33:59	WEB	someone.] ... someone Ok.	
7	1:34:08	INT	Luke knows Isabel. (Inaudible).	A 2
8	1:34:26	WEB	So Sean is the brother of Isabel and she plays basketball.	A 2
9	1:34:32	REC	Sean. So there's a question mark on September 3rd for Sean.	A 3
10	1:34:38	WEB	Ok.	
11	1:34:47	REC	What did? What was he doing?	
12	1:34:50	WEB	He swims 4 to 5 at the outdoor swimming pool.	
13	1:35:19	REC	That's really close to the Career center though I just realized that.	A 4
14	1:35:22	INT	You said Isabel made a deposit?	
15	1:35:27	WEB	Made a deposit of, someone paid Isabel $820.	
16	1:35:30	INT	Ok. Luke owed Isabel money.	A 2
17	1:35:30	WEB	How much?	
18	1:35:37	INT	It doesn't say but he owed her money for a while.	
19	1:35:40	WEB	But Luke's not on here though.	A 4
20	1:36:06	INT	Luke doesn't have many connections. I mean, he has connections with Brandon and Greg but neither of them are on there and that would have been any of the crimes after the 3rd because on the 4th they had student election which they were meeting every Friday between 8 and 10 to work on so they were busy at that time.	A 2

Fig. 6. Example of how artifacts were used together by teams. In turn 3, REC references Isabel from A1. WEB and INT refer to (different parts of) A2 in turns 7–8. In turns 9 and 13 REC refers to A3 and then A4. INT subsequently refers back to A2. And so on. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

The team recorded this connection in their prime suspects table A4 in yellow. The discussion then turned to Luke as a suspect, noting that Luke owed money to Isabel, but also noting that Luke has not been included in the prime suspect table A4. The team eventually concluded that Sean was the thief and Isabel the instigator. The video showed them writing over the original claim (recorded in yellow) in black ink "sold via Isabel" (see Fig. 6).

Prior to the episode in Fig. 6, team 2 had developed a variety of artifacts to summarize and analyze aspects of the problem information. The team created elimination tables, A1 and A3, indicating judgments about who had opportunity to commit each of the four crimes. Based on these artifacts, they decided to create the social network diagram A2 to analyze relationships among those POIs who had opportunity. A1, A2 and A3, as described in Fig. 6,

were then used to derive the prime suspects tables A4 the cells of which include a variety of facts and deductions about possible accomplices. This example illustrates the strategies discussed above: The elimination tables use accretion in comprehensively enumerating all POIs, but also filtering, in their focus on opportunity. The relationship diagram is a filtered successor artifact in focusing on connections among POIs who had been judged to have opportunity. The prime suspects table is further filtered to enumerate only the most likely suspects. The interaction depicted in the episode shows how the team coordinated all four artifacts in developing the analysis.

Although this was a good example of sophisticated development and use of artifacts, the team actually made an incorrect analysis, even though all the necessary information was contained in artifacts A2 and A3. In assessing Luke and Isabel's social connections in the relationship diagram A2, they overlooked his connection to Tay ("Tay" is underlined in blue in A2). Tay and Luke had a history of stealing together (the label on their shared link in A2), and Tay was available at the time of the crime (A3). After the Phase II, when the received updated information, they wrote the correct answer in red ink in the prime suspects table (A4), flagging Isabel and Luke as instigators, in the cell corresponding to Tay and crime 2.

11. Interactions with team performance

In order to articulate performance-related phenomena we selected five consistently high performing team and five consistently low performing teams for microanalysis. High performing teams created fewer information artifacts, 13/35 or 37% of the total (mean=2.6 artifacts per team) than did low performing teams (22/35; mean=4.4) across the three phases, Table 5. Thus, the sheer number of artifacts created does not appear to be a marker or enabler of better team performance.

High performing teams also did not differ much from low performing teams with respect to the specific types of artifacts they created; Table 5 summarizes frequencies across categories of artifacts created. About 38% (5/13) of the artifacts created by high performing teams were relatively complex (tables and graphical representations; see Sections 5 and 6); for low performing teams this proportion was 45% (10/22). Thus, the complexity of artifacts created does not appear to be associated with better team performance.

About 31% (4/13) of the artifacts created by high performing teams were used primarily for information accretion, as discussed in Section 9, and 46% (6/13) were primarily used for information filtering; for low performing teams these proportions strongly reversed: 68% (15/22) were accretion, and 23% (5/22) were filtering (for 5 cases it was not clear what the preponderant strategy was). Based on analysis of videos and transcripts, we classified 69% (9/13) of the artifacts created by high-performing teams as being as coordinated with other artifacts (see Section 10); 59% (13/22) of the artifacts created by low-performing teams were analyzed as having been coordinated in use with other artifacts. Both of these

Table 5
Artifact categories and practices for high and low performing teams.

		High performing teams	Low performing teams
Categories	Text annotation	1	1
	Calendar annotation	1	1
	Map/geo-spatial annotation	0	1
	List	6	9
	Table	3	7
	Graphical representation	2	3
Strategies	Accretion strategy	4	15
	Filtering strategy	6	5
	Major strategy not clear	3	2
Coordinated use		9	13
Total		13	22

patterns suggest that high performing teams adopted more efficient and sophisticated strategies, actively filtering information they inscribed in artifacts, and coordinating their use of multiple artifacts, though with the small number of teams that were assessed in detail, this is a just a statistically descriptive result.

In general, teams were quite bottom-up/opportunistic in creating artifacts. They tended to sketch and prototype artifacts designs, without fully planning or specifying in advance how the artifacts would be used. Two teams were exceptions to this (2 and 13), each had one member who was distinctively concerned with problem decomposition from early in the session. These two teams were also the two highest performing teams. This suggests that the basis for high performance might reside in avoiding premature commitment and analogous consequences of insufficient analysis and reflection [29]. Early top-down analysis is also consistent with creating fewer artifacts overall.

Individuals can make a huge different in the performance of 3-person teams. Our single highest performing team (team 2) created well-organized, complex artifacts primarily designed by one member, including all of the artifacts described in Fig. 6. Only one of their artifacts (depicted in Fig. 5) was created through collaborative effort, though, perhaps notably, that artifact was the final artifact they designed and arguably was the most sophisticated design the team produced.

Of course, we cannot know whether the high performing teams were high performing because they created more filtered artifacts that they also effectively coordinated in use, or whether they created more filtered artifacts and coordinated use because they included members with better problem-solving skills (that is, because they were high-performing). It could even be that our high performing teams succeeded more because they (perhaps fortuitously) created better artifacts earlier on in their problem solving, and thus escaped being caught up in

Table 6
Requirements for interactive visual computing tools and environments to support information analysis.

Task orientation and awareness: Users should be incidentally guided and reminded of first-order task structures and strategies and of resources, contributions and activity of their teammates by the visual presentation of and interaction with tools and materials

Individual annotations: Users should be able to annotate problem documents directly, and easily create (temporary) personal artifacts (such as simple lists)

Individual and team views: Users should be able to share annotations and other artifacts with team members, but should have control over what is shared

Grounding analyses: Users should be able to reinstate the original context of an annotation (that is, they should be able to see the entire object, e.g., document, part of which was annotated)

Successor artifacts: Users should be able to directly edit and further refine an artifact into a successor artifact

Coordinating multiple artifacts: Users should be able to refer to multiple (complementary) artifacts simultaneously

Artifact operations: Users should be able to combine artifact features into filtered views and custom views (they should be able to filter annotations in a map by their date attribute, embed a map in a calendar day view cell, etc)

artifact thrashing, wasting their time creating too many unintegrated artifacts, and thereby eventually experiencing poorer outcomes.

12. Implications for design

Complex cognitive and collaborative activities are often achieved through artifacts and artifact-centered practices. Chin et al. [14] study of five information analysts emphasized their use of annotation and representation in organizing their data to search for patterns, their strongly held attitudes about the diversity of representations used, and the unmet need for information analysis tool support. We found that participants created a variety of artifacts that they used as external cognitive aids. Individual artifacts tended to be simple in structure, mostly direct annotations on problem documents and lists. Teams also created and used artifacts to support collective cognition, that is, to orient and mediate problem solving interactions among team members. Team artifacts tended to be somewhat more refined and complex, for example, including tables and graphical representations.

We analyzed the kinds of artifacts participants spontaneously created, and the ways they used these artifacts. We found that both individual and team artifacts often originate in the spontaneous and iterative tinkering of a single team member. Nevertheless, the team artifacts sometimes become quite articulated, and were used as special purpose tools, closely coordinated with other artifacts (as in Fig. 6). We distinguished information accretion and information filtering in artifact practices, and design of artifacts as successors to prior artifacts. Although our study was qualitative, we also found suggestions that better integrated artifacts and artifact practices were associated with better team performance in our model task.

Our primary purpose in this study was to observe and analyze team performance in information analysis activity supported by non-interactive technologies in order to identify requirements and design ideas for collaborative systems to support such activity. In Table 6, we identify seven design requirements for interactive visual computing tools and environments suggested by this analysis of the use of cognitive artifacts in information analysis activity. These

seven requirements are a design-oriented view of our results; we discuss each in turn.

Information analysis activity is complex, collaborative problem solving. It involves developing and maintaining an orienting model of the task (*task orientation*; [11]), and awareness of team status and strategy with respect to the activity (*activity awareness*; [10]). The information space our participants analyzed – 222 propositions – was modest relative to the real-world, but large enough to challenge teams, particularly given the time pressure in our study. In our problem context, task orientation includes conceptualizing an overall event timeline and geo-spatial frame, and identifying and systematically pursuing strategies such as successive elimination of POIs from consideration. Our teams all did this to some extent, but even the most successful teams directed little specific effort toward such meta-cognitive objectives. In our problem context, activity awareness is keeping track of active team strategies, of what each member can contribute (with respect to role information, but also with respect to insights and perspectives shared in the teamwork to that point), and of what each member is doing currently with respect to active strategies and relevant problem information. Again, our teams did this to some extent, but they did not direct effort to it specifically, and they often did it poorly. (See also [3].)

Human problem solvers characteristically act before they have fully analyzed a problematic situation; indeed, professional problem solving involves a constant self-monitoring to achieve reflection-in-action (e.g., [29]). Getting adults to be more meta-cognitive by directly instructing them to do so is not an effective design intervention (e.g., [11]). A more promising direction, and one quite typical of HCI design work, is to configure interactive environments so that users incidentally encounter and are continuously reminded of high-order problem information in the course of thinking about and acting on immediate and lower level issues. Thus, structuring an interactive environment to build an index of user annotations with respect to POIs, place names, dates, etc. that can be viewed as a list or table, and to provide timeline and map widgets which can be populated with event information that has been annotated might help users develop and maintain a more comprehensive orienting model of the information analysis task. Analogously, keeping track of which team member (that is,

which role) created a given annotation (or other artifact), for example by color-coding, would support activity awareness. For an interactive system to support distributed collaborations, color-coding could also be employed to organize buddy lists and chats. The more pervasively we can exploit such a mechanism, the more effectively we can maintain awareness of what partners know and can contribute to teamwork, and what they have already contributed and currently doing with respect to collective strategies.

Our participants spent substantial time and effort creating a wide variety of artifacts to externalize and support their team cognition. Although Chin et al. [14] study of professional information analysis was smaller scale and less analytic with specific regard to the creation and use of artifacts, they indicate also that artifacts were spontaneously created. This suggests that interactive tool support for information analysis activity should provide templates for cognitive artifacts, or toolkits to create such artifacts. This requirement could of course be realized in many different ways. For example, users could annotate problem documents directly, writing an annotation layer onto document objects, and extracting key content from the annotated data (person names, date, places). Users could be provided with editors to create simple lists and tables.

Although our study data did not capture much detail regarding individual artifact processes, all but one of our 30 participants created individual artifacts, an average of almost 4 per person in the experimental session. Moreover, we observed that team artifacts often originate in the initiative of a single team member. This suggests that supporting easy personal artifact creation/instantiation, a wide range of artifact types, and a mechanism to share artifacts with team members would support and encourage the activities we observed, and perhaps improve them by making sharing an explicit decision. In prior work on collaborative support for emergency management planning, we found that supporting an articulation of public and private views, and making sharing an explicit action facilitated developing an understanding of what other team members knew, and of what information would be useful to share with them; this was manifest, for example, in a strategy shift from asking fellow team members to share information to recognizing what partners were interested in and pushing that specific information to them, a relatively high level of activity awareness among team members [18].

We think that the relatively large number of personal artifacts created in our information analysis activity is a potential resource to the team-level activity, and that its utility could be enhanced if the interactive environment can effectively support and encourage sharing of personal artifacts. One way to realize this is through implementing personal and team views of the problem space, as was done by Convertino et al. [18] in the emergency management planning domain.

We identified a tension between generalization and grounding. In our study, and in Chin et al. [14] prior study, relatively refined and better integrated artifacts that filtered information and provided specialized views were useful to teams, and associated with better performance.

However, it was also important to be able to track back from an artifact to the raw data it codified. This requirement is particularly difficult to support with non-interactive, paper-based artifacts, as in our study and that of Chin et al. [14]. However, linking data representations is one of the primary strengths of interactive systems. We think would be useful to maintain a pointer to original problem document context and allow users to directly access that context from any fact inscribed in an information artifact. For example, the user could select "Holly works @ Starbucks" in Fig. 2A, or the blue line segment labeled REC in the column labeled THURS in the graphical artifact in Fig. 5, and display grounding (presumably a sentence in one or more problem documents).

This functionality could be facilitated and elaborated by allowing users to manually add metadata to artifacts as they are created; for example, indicating social or intentional context beyond what literally appeared in the problem documents. Thus, a user could record that a fact was added to an artifact because of a specific suspicion about a given POI, or over an objection of one team member, or that a given fact in a given artifact is related to another fact inscribed in another artifact, etc. (We acknowledge that creating manual metadata is an example of the kind of meta-cognitive activity people rarely engage in.)

We observed attempts to evolve artifacts, sometimes by merely adding data within a framework, but sometimes evolving the representational framework of the artifact. This latter pattern of evolving successor artifacts was not typical among our teams, but we feel it is significant because it illustrates creativity and critical thinking. Furthermore, evolving artifact designs is not easy to do with physically codified artifacts. One can scribble in an extra data dimension on a paper artifact (e.g., Fig. 3), but this increases clutter and makes the design of the artifact more chaotic. One cannot reuse, adapt, or fundamentally refactor paper artifacts in the way one easily can digital artifacts. Therefore, we feel that the context of an interactive environment for supporting information analysis also provides a much better opportunity to support significant artifact redesign.

In several quite successful episodes in our study, we observed teams coordinating multiple specialized artifacts (Fig. 6). Indeed, were we to better support the creation of relatively refined and better-integrated artifacts that filter data and provide specialized views of the problem data, we might expect to see more interest in coordinating multiple specialized artifacts. Again, this is a particular strength of interactive systems, where multiple coordinated views of an underlying data structure or process are often used to facilitate comprehension and design [7]. Computational approaches to supporting collaborative information analysis have investigated specific utilities of representing underlying data entities, which may be visualized in multiple ways [2,30].

To some extent, our participants needed to create specialized artifacts because the artifacts they could create – out of paper, post-its, and other basic physical materials – were not capable of interaction. Thus, in Fig. 6 there are two kinds of tables being used by the team, but they have

the same axes. In an interactive system context, one might consider that these artifacts could be different views or states of a single interactive artifact. Similarly, we mentioned that the graphical representation in Fig. 4 was derived from a prior artifact, a table. In an interactive system context, we might consider that the table and the graphical representation could be two views of a single underlying data structure. Thus, besides supporting easier creation of specialized artifacts per se, we think it would be useful to provide artifacts with editable views, allowing any given information artifact to a have variety of specialized views.

Drawing implications for design directly from the study of human activities is not straightforward. Key characteristics of human activity co-evolve with the technologies, tools and environments that support those activities [12]. Historically, it is typical that the needs and preferences that motivate a technological innovation are not satisfied tout court by that innovation, but rather transformed into new needs and preferences that are just as keenly felt (and which subsequently motivate further technology development). We are also mindful of those notable cases where the best of intentions have led technologists to over-support apparent inefficiency that was in fact well-adapted vagueness, to no one's benefit [32]. This is not to argue against taking human activity seriously in informing technology design; it is just to put it in the larger and sobering context of history.

More generally, identifying requirements per se is always highly under-constrained. In practice, when requirements are specified and realized in software and in user interaction, given sets of requirements are found to be mutually inconsistent or contradictory, and/or to allow, imply, or even necessitate further requirements. This is the basis for Brooks' well-known argument for the necessity of prototyping ([4]; cf. [8]). Again, this is not to impugn requirements identification and analysis as a key activity in upstream design, it is just to put it in context. The requirements identified constitute a rich design agenda for supporting information analysis with interactive tools. We are currently prototyping embodiments of these requirements. Our approach to this is to selectively respond to, realize, and assess requirements enumerated above, not assuming that in the end we will implement everything, or necessarily implement anything just as envisioned above.

Acknowledgments

References

[1] C.M. Alifano, Fundamentals of Criminal Investigation, Worldwide Law Enforcement Consulting Group, Inc., New York, 2006.

[2] E.A. Bier, S.K., Card J.W. Bodnar, Entity-based collaboration tools for intelligence analysis, in: Proceedings of the IEEE Symposium on Visual Analytics Science and Technology, 2008 (VAST '08), 2008, pp. 99–106.

[3] M. Borge, C.H. Ganoe, S. Shih, J.M. Carroll, Patterns of team processes and breakdowns in information analysis tasks, in: Proceedings of ACM CSCW Conference, Feb 11–15, Seattle, WA, ACM Press 2012.

[4] F.P. Brooks, The Mythical Man-month, Addison-Wesley, Reading, MA, 1975.

[5] J.M. Carroll, What's in a Name?, Freeman, New York, 1985.

[6] J.M. Carroll, H.M. Mentis, G. Convertino, M.B. Rosson, H. Sinha, D. Zhao, Prototyping collaborative geospatial emergency planning, in: B. Van de Walle, P. Burghardt, C. Nieuwenhuis (Eds), Proceedings of ISCRAM 2007: Intelligent Human Computer Systems for Crisis Response and Management, Delft, the Netherlands, May 13–16, 2007, pp. 105–113. ⟨http://www.iscramlive.org⟩.

[7] J.M. Carroll, M.B. Rosson, Deliberated evolution, Human–Computer Interaction 6 (1991) 281–318.

[8] J.M. Carroll, M.B. Rosson, G. Chin, J. Koenemann, Requirements development in scenario-based design, IEEE Transactions on Software Engineering 24 (12) (1998) 1156–1170.

[9] J.M. Carroll, S-I. Shih, C.H. Ganoe, A reference task for collaborative information analysis and decision-making, submitted for publication.

[10] J.M. Carroll, M.B. Rosson, G. Convertino, C.H. Ganoe, Awareness and teamwork in computer-supported collaborations, Interacting with Computers 18 (1) (2006) 21–46.

[11] J.M. Carroll, The Nurnberg Funnel: Designing Minimalist Instruction for Practical Computer Skill, MIT Press, Cambridge, MA, USA, 1990.

[12] J.M. Carroll, W.A. Kellogg, M.B. Rosson, The task–artifact cycle, in: J.M. Carroll (Ed.), Designing Interaction: Psychology at Human–Computer Interface, Cambridge University Press, New York, NY, USA, 1991, pp. 74–102.

[13] J.M. Chaiken, P.W. Greenwood, J.R. Petersilia, The Criminal Investigation Process: A Summary Report, RAND Corporation, Santa Monica, CA, 1976.

[14] G. Chin, O.A. Kuchar, K.E. Wolf, Exploring the analytical processes of intelligence analysts, in: Proceedings of the 27th International Conference on Human Factors in Computing Systems, ACM (2009), 2009, pp. 11–20.

[15] E.G. Chrysikou, R.W. Weisberg, Following the wrong footsteps: fixation effects of pictorial examples in a design problem-solving task, Journal of Experimental Psychology: Learning, Memory, and Cognition 31 (5) (2005) 1134–1148.

[16] M. Cole, Cultural Psychology: A Once and Future Discipline, The Belknap Press of Harvard University Press, Cambridge, 1996.

[17] A. Collins, W. Ferguson, Epistemic forms and epistemic games: structures and strategies to guide inquiry, Educational Psychologist 28 (1) (1993) 25–42.

[18] G. Convertino, H.M. Mentis, A. Slavkovic, M.B. Rosson, J.M. Carroll, Supporting knowledge sharing and awareness in distributed emergency management planning, ACM Transactions on Computer–Human Interaction 18 (4) (2010). Article 22.

[19] R. Cox, Representation construction, externalized cognition and individual differences, Learning and Instruction 9 (1999) 343–363.

[20] P. Dourish, Where the Action is, MIT Press, Cambridge, MA, 2001.

[21] D. Gentner, J. Nielsen, The Anti-Mac interface, Communications of the ACM 39 (8) (1996) 70–82.

[22] E. Hutchins, Cognition in the Wild, MIT Press, Cambridge, MA, 1995.

[23] B. Jordan, A. Henderson, Interaction analysis: foundations and practice, Journal of the Learning Sciences 4 (1) (1995) 39–103. (Maxfield).

[24] J. Larkin, H. Simon, Why a diagram is (sometimes) worth ten thousand words, Cognitive Science 11 (1987) 65–99.

[25] D.A. Norman, Cognitive artifacts, in: John M. Carroll (Ed.), Designing Interaction: Psychology at the Human–Computer Interface, Cambridge University Press, 1991, pp. 17–38.

[26] E. Rosch, Principles of categorization, in: E. Rosch, B.B. Lloyd (Eds.), Cognition and Categorization, Lawrence Erlbaum Associates, Hillsdale, New Jersey, 1978, pp. 27–48.

[27] Rosettex Technology, BearingPoint, & Virginia Tech, Center for Human–Computer Interaction. Collaborative Production Project—Final Merged Report Prepared for the National Technology Alliance, National Imagery and Mapping Agency, Publication Number: TR-001-093003-016, 2004.

[28] W. Schafer, J.M. Carroll, S. Haynes, S. Abrams, Emergency management planning as collaborative community work, Journal of Homeland Security and Emergency Management 5 (1) (2008) 1–17. (Article 10).

[29] D. Schon, The Reflective Practitioner, Basic Books, New York, 1983.

[30] J. Stasko, C. Görg, Z. Liu, Jigsaw: supporting investigative analysis through interactive visualization, Information Visualization 7 (2) (2008) 118–132.

[31] K. Stenning, J. Oberlander, A cognitive theory of graphical and linguistic reasoning: logic and implementation, Cognitive Science 19 (1) (1995) 97–140.

[32] L.A. Suchman, Do categories have politics? The language-action perspective reconsidered, Computer Supported Cooperative Work 2 (1994) 177–190.

[33] L.S. Vygotsky, Interaction between learning and development, in: M. Cole, V. John- Steiner, S. Scribner, E. Souberman (Eds.), Mind in Society: The Development of Higher Psychological Processes, Harvard University Press, Cambridge, MA, 1978, pp. 79–91.

[34] N. Warner, L. Burkman, C.H. Biron. Special Operations Reconnaissance (SOR) Scenario, Intelligence Analysis and Mission Planning, No. NAWCADPAX/TM-2008/184, 2008.

[35] S. Whittaker, L. Terveen, B. Nardi, Let's stop pushing the envelope and start addressing it: a reference task agenda for HCI, Human–Computer Interaction 15 (2000) 75–106.

[36] Wright, J. McCarthy, Technology as Experience, MIT Press Cambridge, MA, 2004.