Factors Contributing to Performance for Cyber Security Forensic Analysis

Shelby Hopkins, Andrew Wilson, Austin Silva, and Chris Forsythe(✉)

Sandia National Laboratories, Albuquerque, NM, USA
shopkinl@trinity.edu,
{atwilso,aussilv,jcforsy}@sandia.gov

Abstract. Previously, the current authors (Hopkins et al. 2015) described research in which subjects provided a tool that facilitated their construction of a narrative account of events performed better in conducting cyber security forensic analysis. The narrative tool offered several distinct features. In the current paper, an analysis is reported that considered which features of the tool contributed to superior performance. This analysis revealed two features that accounted for a statistically significant portion of the variance in performance. The first feature provided a mechanism for subjects to identify suspected perpetrators of the crimes and their motives. The second feature involved the ability to create an annotated visuospatial diagram of clues regarding the crimes and their relationships to one another. Based on these results, guidance may be provided for the development of software tools meant to aid cyber security professionals in conducting forensic analysis.

Keywords: Cyber security · Forensic analysis · Decision making · Narratives

1 Introduction

On a weekly basis, reports appear in the media describing a data breach, denial of service or other cyber crimes involving major corporations, governments or military organizations. As those who perpetrate cyber crimes become increasingly sophisticated, there is a growing gap in the number of available cyber security analysts and the demand for their services (Burning Glass 2014). Solutions are needed to enhance the effectiveness of current cyber security analysts while accelerating training of those entering the field.

Within many organizations, cyber security analysts conduct activities comparable to criminal forensic analysis. The analysts piece together clues to understand a series of events, including the likely perpetrator and their objectives and capabilities. The current researchers have experimentally demonstrated that mechanisms facilitating and encouraging construction of a narrative account of events produced superior performance in a cyber forensic analysis task (Hopkins et al. 2015). Research is reported here that considered which features of a narrative representation of events accounted for superior outcomes.

T. Tryfonas and I. Askoxylakis (Eds.): HAS 2015, LNCS 9190, pp. 200–206, 2015.
DOI: 10.1007/978-3-319-20376-8_18

2 Methods

2.1 Subjects

Subjects consisted of 26 employees of Sandia National Laboratories who responded to a company-wide announcement soliciting volunteers to participate in a research study concerning criminal forensic analysis.

2.2 Materials

A scenario was composed based on publicized reports of cyber crimes. The scenario involved a fictitious pharmaceutical manufacturer and subjects were given the pretense that they had been asked to investigate a series of suspicious events at this company. The scenario involved three separate crimes committed by three distinct entities operating independently of one another and with different motives and objectives. The first scenario involved a Hacktivist group intent on proving the pharmaceutical company was involved in controversial activities (i.e., biological weapons research). In the second scenario, a criminal organization committed bank fraud with funds stolen from accounts used by the company. The third scenario consisted of intellectual property theft by an employee of the company (i.e., Insider).

For each crime, a collection of clues were created that realistically, would be available to a corporate security officer conducting a forensic analysis. There were a total of 16 legitimate clues with the Hacktivist thread being the more complex having 8 clues, and the Criminal and Insider threads being somewhat simpler with 4 clues each. There were eight additional clues that served as "red herrings" and had nothing to do with the three crimes. Laminated cards presented a one sentence description of the clues and the associated date the event was noted. Two cyber forensic analysts reviewed each scenario and verified that the storyline and clues were plausible and representative of the types of crimes a cyber forensic analyst might actually encounter.

2.3 Procedure

Subjects were randomly assigned to either a Narrative or Association condition.

Narrative Condition. Subjects were provided 24 laminated cards with magnetic backings on which the clues and associated dates were printed. Subjects were asked to work at a 57" × 46" magnetic whiteboard. Subjects arranged the clues by affixing them to the whiteboard, and used dry erase markers (black, blue, green and red) to draw links between clues and boundaries encircling groups of clues, as well as make notes and other markings. As shown in Fig. 1, features were provided to facilitate and encourage subjects to construct a narrative. Narrative features included 5 Criminal Entity Cards with labeled spaces for subjects to use dry erase markers to denote the identity of the entities, "What trying to do?" and "Why trying to do it?," and a timeline spanning a period encompassing the dates associated with the clues. The upper right corner of the board was labeled "Red Herrings" to encourage subjects to segregate legitimate and red herring clues and subjects were given 12 annotation cards on which to make notes,

8 context cards to identify contexts, and circular magnets to use as tags with 5 different colors (white, blue, green, yellow, and red) and 6 magnets in each color (total of 30 magnets). The board also had a vertical axis labeled, "Criminal Entities," and a horizontal axis for the timeline with months of the year denoted as tick marks. Once subjects had indicated they understood the assignment, they were given a box with the clues arranged in a random order and allowed 25 min to conduct their analysis. At the conclusion of the test session, a photograph was taken of the diagram created by the subject for subsequent data analysis.

Association Condition. The Association condition provided the same visuospatial elements as the Narrative condition, but without features to facilitate construction of a narrative. The same laminated cards with clues were provided and work was completed at the whiteboard. However, subjects were only provided with dry erase markers and the colored circular magnets. Subjects were instructed that the goal of this task was to identify clues that were related to one another and then, signify any relationships between the groupings of clues using the dry erase markers or colored magnets. Figure 2 shows the diagram created by one of these subjects.

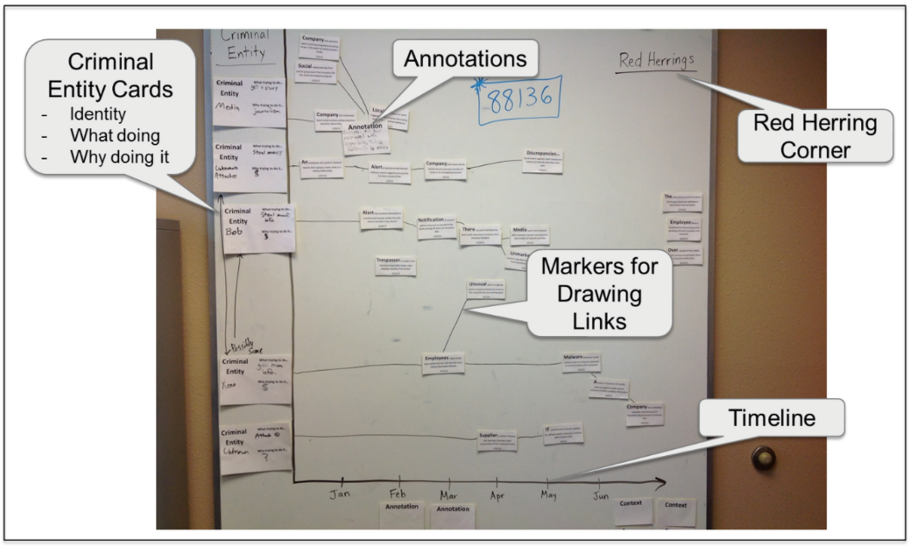

Fig. 1. Example of the whiteboard configuration and features provided to subjects in the narrative condition. Magnetic markers that could be used as tags are not shown here (Color figure online).

Following the forensic analysis, subjects were asked to depict their interpretation of the events using the software tool PlotWeaver (See Fig. 3). PlotWeaver provides an XML-based graphical interface for creating pictorial representations of events. In diagraming stories, PlotWeaver allows entities and interactions between entities to be identified as a time-dependent series of events. Subjects were provided a brief tutorial

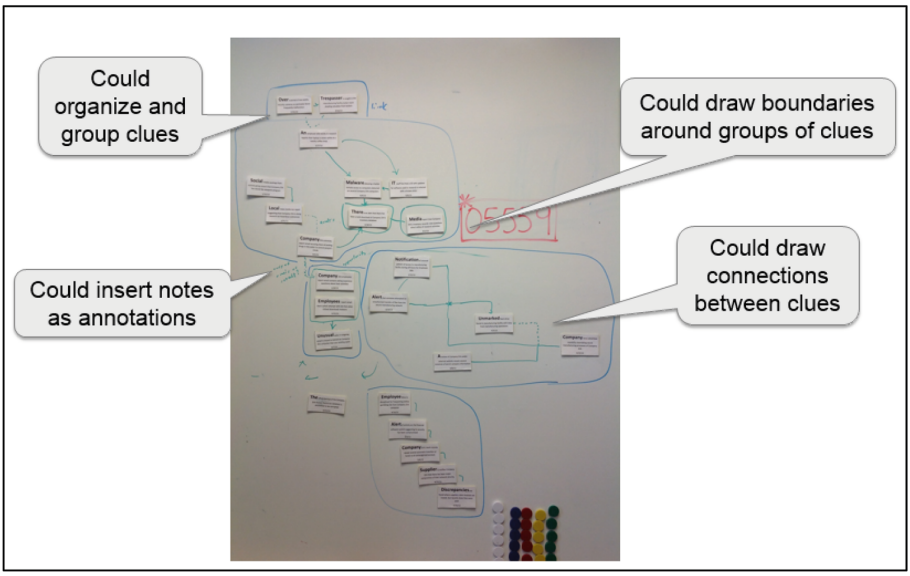

Fig. 2. Example of diagrams created by subjects in the association condition (Color figure online).

on how to use the key features of PlotWeaver. Once the experimenter had verified that subjects understood these features, subjects were given 25 min to create their PlotWeaver interpretation of events. During this time, whiteboard diagrams created by subjects were available and could be referenced at any time.

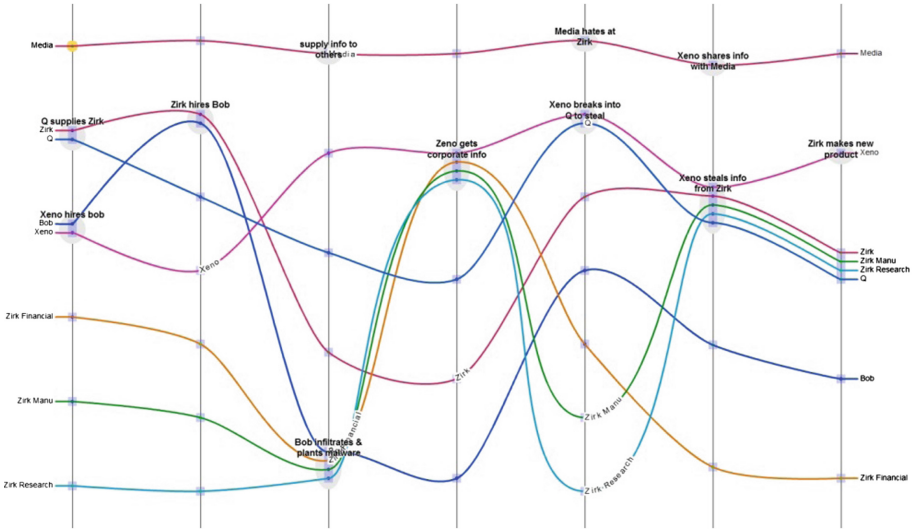

Fig. 3. Example of a PlotWeaver diagram illustrating the subject's interpretation of events within the scenario.

3 Results

Within the diagrams, there were twelve distinguishable features identified that were available to subjects in both the Narrative and Association conditions (See Table 1). The initial analysis considered the relative usage of these features by subjects in each condition. The average frequency that ten of the twelve elements appeared in the whiteboard diagrams of subjects is presented in Fig. 3. It may be observed that on average, subjects in the Narrative condition identified more entities ($t = 5.50$; $p < 0.001$) and more entity motives ($t = 7.48$; $p < 0.001$) than those in the Association condition. For the other features shown in Fig. 3, the differences in usage between the two groups were not statistically significant (Fig. 4).

Table 1. Features identified that were available in constructing diagrams for the narrative and association conditions.

– Group
– Group label
– Group annotation
– Links
– Link labels
– Entities
– Entity motives
– Tag categories
– Tags used
– Questions
– Chronological order
– Red herring segregation

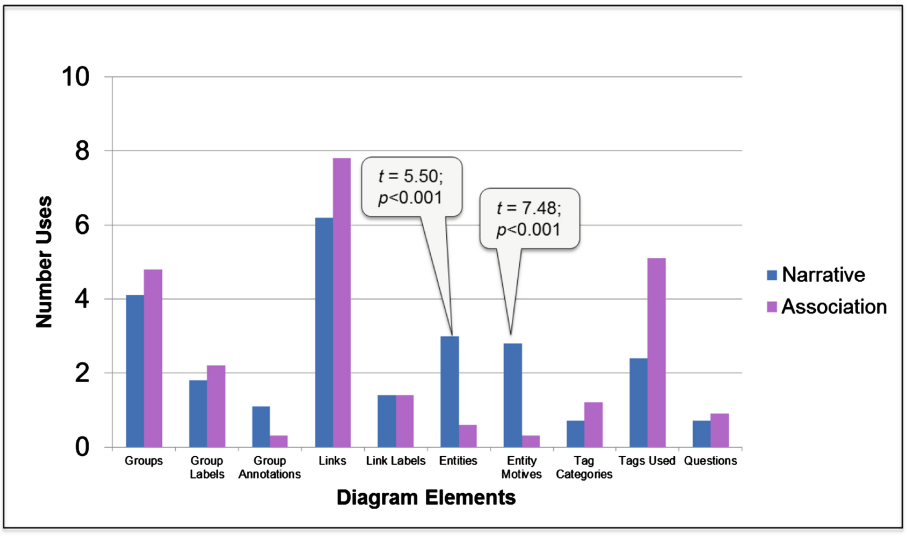

Fig. 4. Average frequency different elements appeared in the whiteboard diagrams of subjects in the narrative and association conditions.

Diagrams were analyzed to identify if subjects had arranged clues in a chronological order. This involved any grouping of three or more clues in which the clues were arranged either vertically or horizontally from earlier to later dates. Groupings may or may not have aligned with the timeline provided in the Narrative condition. It was found that subjects in the Narrative condition were more likely to use chronological orderings of clues ($t = 3.99$; $p < 0.001$).

To determine if subjects had segregated red herring clues from other clues, arrangements of one or more clues were identified where there was no connection made to any of the other clues in the diagram (e.g., no lines were drawn linking clues to other clues or providing the boundary for a group of clues). It was found that subjects in the Narrative condition were more likely to segregate clues that were unassociated with the other clues (i.e., red herrings) ($t = 3.16$; $p < 0.001$).

Groups were compared with regard to how many different types of features and the total number of elements used in the diagrams. It was found that the subjects in the narrative condition used more different types of elements than those in the Association condition ($t = 3.27$; $p < 0.01$). However, when the total number of elements appearing in diagrams was considered (e.g., each arrow connecting clues and each circular magnet was counted as a separate element), the groups did not significantly differ. On average, subjects in the narrative condition used 27.1 distinct elements in their diagrams and subjects in the Association condition used 25.8 elements.

From an analysis of the PlotWeaver storylines constructed by subjects to depict their interpretation of events, it was found that subjects in the Narrative condition used more clues within storylines (Hopkins et al. 2015). Furthermore, while subjects in the Narrative condition included more of the legitimate clues in their plots, both groups included approximately the same number of red herring clues. Thus, subjects in the Narrative condition, developed richer storylines while being no more susceptible to the red herrings. Using each of the twelve features analyzed for the whiteboard diagrams, a stepwise regression was calculated to identify which features accounted for the superior performance of the subjects in the Narrative condition. When the number of clues included in the plots was considered, the resulting model was statistically significant ($F = 6.55$; $p < 0.01$) with three factors (Group Annotations, Entities and Entity Motives) accounting for 60 % of the variance in performance ($R^2 = 60.2$, R^2 *adjusted* = 51.0). Similarly, when the number of legitimate clues included in the plots was analyzed, the resulting model was statistically significant ($F = 4.55$; $p < 0.05$) with the same three factors accounting for 51 % of the variance in performance ($R^2 = 51.2$, R^2 *adjusted* = 39.9).

4 Conclusion

Subjects provided with features to facilitate and encourage their construction of a narrative account of events used these features. In particular, they used the entity cards to identify suspected perpetrators of the cyber crimes and their respective motives, while organizing events within a chronological order. The diagrams constructed by subjects in the Association condition that incorporated comparable visuospatial features, but without the narrative components, included an equivalent overall number of

elements. However, the elements were used in a less diverse manner and were less likely to incorporate the features of a narrative. Furthermore, subjects in the Narrative condition performed better with regard to their accurately inferring the nature of the crimes committed, while being no more susceptible to clues that were red herrings.

Further analysis revealed two diagram features that accounted for the performance of the subjects in the narrative condition. First, subjects that focused more on identifying suspected perpetrators and their motives did better than those that devoted less of their analysis to these factors. This suggests that forensic analysts benefit from artifacts that facilitate their ability to discern who may have committed a crime and why they may have done so. Second, subjects that arranged clues in groups and annotated those groups did better than subjects who did not do so. This result may be attributed to the annotations providing memory support in allowing subjects to remember associations and inferences regarding the clues. The scenario used in the current study was complex and somewhat ambiguous. The ability to use annotations may have provided valuable memory support in coping with this complexity.

It is not surprising that the consideration of entities and their respective motives was a primary predictor of performance in the current study. This finding is consistent with event indexing models of narrative comprehension in that entities are considered to be a primary dimension along which stories are organized (Zwaan and Radvansky 1998). Furthermore, entity intentions is a second dimension within this model. It is surprising that the ordering of clues chronologically was not a factor, given that time is also a factor within the event indexing model. The current study did not provide features allowing subjects to organize clues with regard to spatial considerations. It was noted that several subjects did consider spatial factors separating the clues involving outside entities from those involving entities within the fictional corporation and distinguishing different domains from one another (e.g., manufacturing and financial). It is suggested that further benefits may be derived from providing features that facilitate the ability to make spatial distinctions and incorporate these distinctions into the overall representation.

Acknowledgements. Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000. (SAND2014-2123 C)

References

Burning Glass: Job market intelligence: report on the growth of cyber security jobs (2014). http://www.burning-glass.com/media/4187/Burning%20Glass%20Report%20on%20Cybersecurity%20Jobs.pdf

Hopkins, S.E., Silva, A., Wilson, A., Forsythe, C.: Facilitation of forensic analysis using a narrative template. In: Proceedings of the Applied Human Factors and Ergonomics Conference, Las Vegas, NV (2015)

Zwaan, R.A., Radvansky, G.A.: Situation models in language comprehension and memory. Psychol. Bull. **123**(2), 162–185 (1998)