Story Web: Case-Based Reasoning to Enhance Analysts' Anticipatory Thinking

Robert E. Wray, Kelly J. Neville, Bob Bechtel, and Tyree Cowell
Soar Technology
Ann Arbor, MI

A workshop was conducted to evaluate the usefulness of a prototype research and analysis support tool in intelligence analysis. It was hypothesized that the tool, called *Story Web*, would enrich intelligence analysts' anticipatory thinking and help them identify a broader and more varied set of possible futures. Story Web uses analogical reasoning to map analyst work products to a database of relevant, past cases extracted from open-source news articles, case studies, and other sources. Matched content from those sources is presented to the analysts for consideration. Eight intelligence analysts volunteered to participate in two (n=4) or three (n=4) analysis sessions. They were instructed to assess, across sessions, potential cyber threats to the 2019 Canadian federal election and to represent their work using concept-mapping software (IHMC's CMapTools). At the start of Sessions 2 and 3, analysts were provided with Story Web-generated recommendations to consider. Seven of the eight analysts reported that Story Web gave them ideas and suggested lines of inquiry they otherwise may not have considered. Most (five) analysts reported that Story Web enriched the existing possibilities they were developing; three reported that Story Web helped them think outside the box, i.e., divergently. All participants suggested improvements to the user interface (which is not yet fully developed) and most offered suggestions related to improving story representations and supporting human-machine teamwork and trust. The results of this workshop will be used to improve the usefulness and usability of Story Web, the interpretability of its recommendations to analysts, and its integration into and support of analysts' work activities.

Background

Story Web is a software-based technology that uses analogical reasoning to draw connections between an analysts' work and past related events. Analogical reasoning is an important strategy in problem solving but can be difficult for humans when similarities are not perceptually evident or pointed out (Gentner & Markman, 1997; Hoffman et al., 2011; Lewandowsky & Kirsner, 2000). Story Web supplements this reasoning and the analyst's knowledge base on which it depends to enrich anticipatory thinking, including both *evolutionary* and *revolutionary,* i.e., *divergent,* anticipatory thinking variants (Voros, 2006). The intent is to help novice analysts forecast as though they were journeymen and journeyman-level analysts to forecast as though they were experts.

In addition to broadening analysts' thinking by drawing analogies to past cases, the Story Web development effort is exploring the use of story-based representations to improve human-computer cooperation in joint activity. Stories and storytelling are hypothesized to play a central role in human intelligence. In particular, the story is an effective and functionally useful organizational structure for intertwining and integrating knowledge. Research suggests that the structure of an unfolding event is one of the dominant means by which human memory is organized (e.g., Klein, Loftus, & Kihlstrom, 2002; Tversky & Thomson, 1973) and that we rely heavily on event- structured memories and knowledge to support sensemaking, decision making, and planning (Klein, 1993). We speculate that intelligence analysts will find Story Web outputs to be more comprehensible and useful when they are based on and communicated in a format that retains story elements.

In its current version, Story Web takes as its inputs articles about past relevant events. In the case of the Story Web evaluation workshop described below, these were articles about cyber threats to democratic elections. From these articles, a series of triples is extracted. The triples, consisting of subject-predicate-

object sequences, are interconnected into a concept map through their mappings to related concepts in an ontology (derived from the *Suggested Upper Merged Ontology* [SUMO] and extended for this workshop by using the *Unified Cyber Security Ontology* [UCSO] [https://github.com/ucoProject/UCO/tree/master/docs]). The resulting concept maps are referred to as *background stories;* these are the stories Story Web draws on to enrich analysts' anticipatory thinking and analysis work. Source articles can consist of either or both event details and general factual information and thus may be considered as contributing to Story Web's *episodic* and *semantic* memory stores, respectively.

Analysis products that Story Web accepts for processing also take the form of a concept map. The analyst can work directly in concept mapping software or, in future Story Web versions, concept maps might be derived from other analysis formats. As an analyst develops an analysis concept map product, Story Web maps the product's contents to background story representations through the designated ontology. When Story Web finds a match, it suggests the matched background story nodes or links and surrounding content (the span of activation can be varied by the user) to the analyst for consideration. Matches can be exact, between concepts in the same ontological category, or between sets of nodes with similar relational structures. Story Web recommends to the analyst matched concepts (link and node labels) in an order based on the match score (the score feature is under development).

Evaluation Workshop

The Story Web analogical reasoning algorithms were originally developed for use in mission and disaster relief planning (Wray et al., 2016). A workshop was recently conducted to evaluate the use of the algorithms in intelligence analysis. The focus was on determining if and how Story Web's analogical reasoning algorithms could benefit intelligence community objectives to improve anticipatory thinking, divergent thinking, and forecasting. We suggest these benefits will take two primary forms:

- Supplementing analysts' knowledge bases so they have a richer and wider variety of knowledge and possibilities to start with and branch from (one of the top ten analyst needs identified in a 2005 analyst workshop; Badalamente & Greitzer, 2008) and
- Helping analysts overcome biases that stem from latching onto a particular *schema* (knowledge restructures used to support sensemaking and decision making; may also be referred to as *mental models*; Klein, Moon, & Hoffman, 2006) when an analyst could instead be actively considering a range of schemas and schema variations.

In preparation for the evaluation workshop, the Story Web algorithms, which had been developed by a different team, were evaluated and refined; background stories were extracted from articles related to the workshop analysis problem, assessing the cyber attack threat to democratic elections; the ontology used in Story Web was extended by adding to it the UCSO; and a rudimentary user interface was developed. The workshop was conducted in a naturalistic setting to allow analysts to evaluate Story Web in the context of typical work activities. The participation of professional analysts translated into a relatively small number of participants, which precluded the use of statistical data analyses. On the other hand, analyst participants were able to provide experience-informed data from which we could distill valuable collective recommendations and unique insights.

Method

Participants

Eight professional intelligence analysts assigned to the Laboratory for Analytic Sciences (LAS) volunteered to participate in a 1-day workshop held in a university classroom. The group participated in three analysis sessions, each followed by a data collection period. Analysis sessions were 1 hr, 45 min, and 30 min in duration, respectively, and data collection periods were approximately 15 min. Following each data collection period, participants were given an approximately 1-hr break.

Before beginning the first analysis session, analysts were given a set of handouts and asked to participate in a 1-hr orientation. Handouts consisted of a workshop overview that included the Story Web team's contact information; problem statement and task descriptions; analysis guidance; pages 11-12 from the Canadian Communications Security Establishment (CSE) report entitled *Cyber Threats to Canada's Democratic Process,* which provide overviews of Canada's democratic process and cyber threats, respectively; concept map development instructions and guidance (we used *CMapTools* [Institute for Human and Machine Cognition [IHMC], https://cmap.ihmc.us/cmaptools/]); an overview of the data collection plan; and questionnaires to be completed after each session.

As a result of competing obligations, four analysts did not participate in the third analysis session.

Procedure

Analysts were instructed to bring with them to the workshop the government laptop computer assigned to them for their daily work, pre-loaded with the IHMC CMapTools software. When they arrived, they were asked to choose a seat and were given a pack of handouts and a print out of PowerPoint slides describing Story Web and the workshop.

At the commencement of the workshop, members of the research team and the analyst participants introduced themselves. Using an overhead projection of the distributed PowerPoint slides, an evaluation team member gave an overview of the Story Web tool and workshop, presented the contents of each handout, and answered analysts' questions.

After the orientation, the participants began their first one-hour analysis session. For each analysis session, the analysts used research resources available on the public internet to identify and characterize potential cyber threats. They were permitted to use any resources they had available, such as pen and paper, in addition to their laptops. They were asked to start translating their work into a concept map at least 15 min before the end of each session. An evaluation team member notified the group when that time arrived. Four evaluation team members were available throughout the session to help analysts and answer questions.

At the start of Sessions 2 and 3, the analysts were given pens and packs of sticky notes, blue sticky notes at the start of one session and green sticky notes at the start of the other. The analysts were instructed to use them if, as they worked, they noticed or thought of Story Web feedback they would like to provide or include in their questionnaire responses. In particular, they were asked to note when:

- Story Web outputs helped them think of a creative or novel concept or path to add to their concept map,
- Story Web outputs gave them an insight or help them think about the problem in a new way, or
- They noticed a short-coming or thought of a potential improvement to the Story Web tool or its outputs.

At the end of each analysis period, participants were asked to complete a pen-and-paper questionnaire. An abbreviated version was used following Session 1. Session 2 and 3 questionnaire questions were aimed at investigating more thoroughly the topics listed above for sticky-note capture. After each analyst completed the questionnaire, he or she left for a scheduled break.

During the breaks, evaluation team members collected completed questionnaires and sticky notes and saved analysts' concept maps onto a USB drive. They ran Story Web on each computer and submitted to it the analyst's Session 1 concept map. Although Story Web is envisioned as a tool that will run and periodically provide suggestions while the analyst works, it currently operates independently of the analysis product's development.

This version of Story Web features a rudimentary user interface that outputs a separate window for each background story. Each output window contains the nodes and links of the analyst's concept map that matched background story elements along with the matching background story elements. The layered stack

of output windows was left displayed on the analysts' computer displays for them to view when they returned.

As each analyst returned, he or she viewed the displayed Story Web outputs. An evaluation team member helped each analyst orient to the set of output windows. Analysts were asked to continue their analysis work as soon as they had finished reviewing the Story Web output windows to their satisfaction.

After Session 3, participants were interviewed. When each analyst finished his or her analysis work, they either completed the questionnaire or participated in an interview, depending on the availability of the four evaluation team members, or they participated in the interview and then completed the questionnaire.

Interviews were semi-structured and each was conducted by a single evaluation team member. The interview protocol specified asking the analyst to discuss recent changes to his or her concept map; in particular, novel, surprising, or creative additions; enriching detail and elaboration; conservative additions; and structural changes. Analysts were asked to try to recall and describe the thought process that led each recalled change. Prompts included:

- What triggered the addition/change? A particular piece of information, memory, or thought?
- For the addition/change you just described, were Story Web suggestions considered? Why or why not?
- How did the presentation format or timing of Story Web suggestions affect your work developing this particular element/change?

The team member conducting the interview captured the analysts' discussion as well as possible using pen and paper. The handwritten notes were later typed up and collated for data analysis purposes. Participants were thanked and left after finishing their interviews.

Results

Questionnaire, sticky note, and interview data were (1) reviewed to identify support for and against the value of Story Web to anticipatory thinking and (2) categorized into themes across participants. Cross-session changes in concept maps are not reported as we instructed participants to create them only to help them evaluate and generate feedback about the potential value of the tool and how it can be improved for use in intelligence analysis work.

One participant provided discussion-based feedback but did not complete questionnaires or participate in the interview. In addition, multiple analysts did not understand that each Story Web output window contained matches from a different background story and only used the output window displayed on top of the window stack. The results below are thus based on the Session 2 data of seven analysts, the Session 3 data of four analysts, and primarily data associated with use of a single Story-Web output window at the start of each session.

Anticipatory Thinking

All seven responding analysts reported that Story Web outputs contributed to the identification of additional concept map content. Feedback suggested that Story Web was more supportive of developing existing lines of inquiry than identifying new possibilities. Six reported that Story Web contributes to the elaboration of existing lines of analysis and two reported that it does not. Three (including two of the six reporting positively about existing lines of analysis) reported that Story Web contributes to the recognition of new lines of analysis whereas five reported that it does not. One participant responded that Story Web does not contribute to either type of concept map addition; however, in response to a different question, reported that Story Web output prompted in Session 2 the addition of two new nodes to her analysis product (she did not participate in Session 3).

Concept map additions found in or based on Story Web outputs ranged from 0 to 8 per session, which we consider reasonable and promising given the limited duration of Sessions 2 and 3 and the time required to

conduct an analysis (We were originally advised to give analysts a three-day workshop, versus one). One was the most commonly reported number of Story Web-based additions per session reported; two was the second most commonly reported. In summary, evidence supports the notion that Story Web potentially will be of value to analysts by pointing them to useful resources and suggesting information the analyst might not otherwise consider.

Feedback Themes

Five major themes were identified in the analysts' feedback. These are summarized below:

Tool Usefulness. Six of the seven responding analysts provided positive feedback about the value Story Web could offer their work. Analysts stated that the tool helped them to identify new ideas and lines of analysis; one described it as having promise.

Output Content Interpretability. Five of the seven responding analysts commented on the interpretability of the content of Story Web outputs. Their feedback points to a need to continue to smith the process used for deriving story representations from articles.

Output Display Readability and Usability. All seven of the responding analysts criticized or provided suggestions about how to improve the output displays used to present matched background story content. The focus of the technology development leading up to the workshop had been on the matching algorithms, not on the presentation of the information, so this was not surprising. This feedback will be invaluable as development continues and more attention can be placed on the technology's interface with the analyst.

Human-Machine Trust and Teamwork. Five of the seven responding analysts provided suggestions that we classified as calling for improved trust and teamwork. These suggestions could improve Story Web in ways that give analysts trust in its products and allow Story Web to support analysts' research and assessment work in more active and responsive ways.

Concept Map Development and Integration. Whereas the themes presented above involve interactions with Story Web and its outputs to analysts, a fifth theme involved the concept mapping software analysts used to create their analysis product. A concept mapping tool was chosen for the simple reason that Story Web was designed to process concept maps. If the analysts had produced another type of analysis product, the evaluation team would have converted it into a concept map before submitting it to Story Web.

We did not know how well the concept mapping tool would be received by the analysts and were surprised to find it relatively well accepted. No analyst rejected or complained about the concept mapping tool and all analysts indicated they would be willing to use it as one their work tools and possibly as a primary work tool. Analysts' suggestions included additional concept mapping tool features and, central among their suggestions, the integration of Story Web with their concept mapping work.

Conclusions

In summary, workshop results suggest the Story Web tool holds value for intelligence analysts by helping them identify more lines of analysis and more details and possibilities within those lines. There was also evidence to suggest that Story Web may contribute to revolutionary, or divergent, anticipatory thinking by helping analysts think more broadly and consider possibilities outside the boundaries of their own knowledge and resource base. According to the feedback provided, Story Web can be shaped into a useful research and analysis support tool for intelligence analysis. Results of the workshop included feedback about specific areas of needed improvement and specific improvements.

Future Story Web development needs to focus on a number of aspects of the background story representation. It needs to focus on providing flexible, searchable access to a wide range of background story sources. In addition, significant feedback called for improving the story representations derived from those sources. In particular, analysts requested improved interpretability of the story representations. Further, multiple analysts indicated they require embedded details or meta-data about the story sources and

content and would also like to be able to directly access the article or article section from which a given match and recommendation is derived.

Secondly, the development team needs to use workshop feedback to design effective Story Web output displays. More generally, it needs to improve and potentially revamp the user interface, which was intended to be temporary, and its ability to support team-like interactivity between Story Web and analysts. Story Web needs a user interface that supports analysts in viewing, navigating, and working with Story Web and its outputs. Ideally, the interface design will support real-time integration between the analyst's work and Story Web's recommendations. This sophisticated interface will require numerous iterations as well as significant teamwork among analysts, an interface designer, and the developer of the Story Web algorithms that determine story representations and match scores.

Acknowledgements

We are grateful to our LAS project leads, Adam Amos-Binks (now with Applied Research Associates) and Sarah-Margaret Tulloss, and to our other LAS collaborators, John Slankas, Colleen Stacy, Lori Wachter, and Jamie Roseborough, for their significant roles in planning and conducting the Story Web workshop. We are likewise indebted to the eight analysts who volunteered their time to evaluate Story Web and generously shared their expertise and feedback. The workshop was a rare opportunity to obtain feedback from core stakeholders and it has given us a much better understanding of analysts' work, the demands they face, and how technology can facilitate that work.

References

Communications Security Establishment (2018). *Cyber threats to Canada's democratic process.* Ottawa, CA: CSE.

Badalamente, R.V., & Greitzer, F.L. (2008). Top ten needs for intelligence analysis tool development. *Proceedings of the 2005 International Conference on Intelligence Analysis* (McLean, VA), https://analysis. mitre.org/proceedings/index.html.

Gentner, D., & Markman, A. B. (1997). Structure mapping in analogy and similarity. *American Psychologist*, *52*(1), 45-56.

Hoffman, R. R., Deal, S. V., Potter, S., & Roth, E. M. (2010). The practitioner's cycles, part 2: Solving envisioned world problems. *IEEE Intelligent Systems*, *25*(3), 6-11.

Hoffman, R., Henderson, S., Moon, B., Moore, D. T., & Litman, J. A. (2011). Reasoning difficulty in analytical activity. *Theoretical Issues in Ergonomics Science*, *12*(3), 225-240.

Klein, G. A. (1993). A recognition-primed decision (RPD) model of rapid decision making. In G. A. Klein, J. Orasanu, R. Calderwood, & C. E. Zsambok (Eds.), *Decision making in action: Models and methods* (pp. 138–147). Norwood, NJ: Ablex.

Klein, G., Moon, B., & Hoffman, R. R. (2006). Making sense of sensemaking 2: A macrocognitive model. *IEEE Intelligent Systems*, *21*(5), 88-92.

Klein, S. B., Loftus, J., & Kihlstrom, J. F. (2002). Memory and temporal experience: The effects of episodic memory loss on an amnesic patient's ability to remember the past and imagine the future. *Social Cognition*, *20*(5), 353-379.

Lewandowsky, S., & Kirsner, K. (2000). Knowledge partitioning: Context-dependent use of expertise. *Memory & Cognition*, *28*(2), 295-305.

Tulving, E., & Thomson, D. M. (1973). Encoding specificity and retrieval processes in episodic memory. *Psychological review, 80*(5), 352-373.

Voros, J. (2006). Introducing a classification framework for prospective methods. *Foresight*, *8*(2), 43-56.

Wray, R. E., Folsom-Kovarik, J. T., Schmorrow, D., Jones, R. M., & Marinier, R. (2016). Using stories to deepen shared human-computer understanding of PMESII systems. In *Advances in Cross-Cultural Decision Making* (pp. 29-41). Springer.