The SAnTA Recommender System and Naturalistic Decision Making

Frank LINTON[b], Sarah BEEBE[a], Mark BROWN[b], Casey FALK[b], and Mark ZIMMERMANN[b]
[a]*Ascendant Analytics*
[b]*The MITRE Corporation*

ABSTRACT

This paper presents a practice-oriented original contribution that advances the application of NDM. SAnTA is a prototype recommender system that is currently being piloted. SAnTA embodies expert reasoning and supports individual and team performance in a complex real-world setting. SAnTA's users are analysts who create reports for decision-makers; the reports are based on data and information from multiple sources. Both the data and the decision-makers' needs (the analysts' sources and goals, respectively) may be dynamic, uncertain and continually changing, consequently analysts must also make decisions about their work processes and products under these same circumstances. SAnTA aids these analysts' decisions by helping them select structured analytic techniques to reduce cognitive bias. Analysts' reports are usually not needed in real-time, but they may be urgent, and mistakes can have significant consequences.

KEYWORDS

Decision Making; Education and training; Externalized/Embedded Cognition; recommender system; structured analytic technique; intelligence analysis.

INTRODUCTION

Unlike much NDM research (Klein, 2008), this work does not address decisions made while fighting a fire or a battle, but instead addresses decisions made when creating the analysis that may prevent, or lead to, a battle – or a war. The analyst's life generally is not on the line, but their professional reputation may be. In this environment, the aspects of NDM that apply are making "…tough decisions under difficult conditions such as limited time, uncertainty, high stakes, vague goals, and unstable conditions." (Klein, 2008)

Structured analytic techniques (SATs), including qualitative aids aimed at promoting critical thinking and mitigating cognitive mind-sets and biases, are increasingly used by analysts and are mandated by analytic production organizations. SATs are now taught in analytic training programs world-wide. SAT usage is supported by reference materials and by facilitators who guide analysts using a question-based approach to select relevant SATs, and apply them in the resulting analytic process. Yet barriers remain: There are relatively few experienced tradecraft analytic advisors, the number of SATs is growing, and the reference materials offer little guidance on how to select the best SAT for the analytic challenge at hand. Thus, the conditions for improper selection and misapplication of SATs are ripe.

SAnTA is a recommender system intended to prompt critical thought, suggest applicable SATs, and provide further resources and guidance on SATs. SAnTA employs a question-based critical thinking approach to prompt consideration of both process- and problem-related factors that affect relevant SAT selection. The tool provides recommendations on the basis of analysts' characterizations of their progress through the analytic process. The resulting recommendations are ranked, and include explanations of the basis for the recommendations and descriptions of the techniques.

SAnTA's end users are essentially characterizing their situation in order to learn what sorts of bias are most likely to be present at this point in the analysis and with which SATs (if any) can these best be mitigated. SAnTA's end users are typically less-experienced, and do not have enough expertise to make optimal SAT selections by recognition-primed decisions (RPD) (Klein, 2008).

Recommender systems are frequently found in web-based commerce, where they recommend purchases to consumers. In that domain, much research is focused on algorithms for generating good-enough recommendations quickly. In contrast, very little recommender research has focused on recommendations as decision aids, helping people make the best choices about what to do next, as SAnTA does (Felfernig, Chen & Mandl, 2011).

Recommender systems are characterized by *how* their recommendations are generated (Ricci, Rokach, Shapira, & Kantor, 2011). They are characterized as collaborative filtering (based on user ratings), content-based filtering (based on item attributes), and knowledge-based filtering (explicit knowledge of how certain items meet users' needs). SAnTA is a recommender of this third type, a knowledge-based recommender. In SAnTA, expert

knowledge pertaining to the topic of selecting and applying structured analytic techniques (SATs) has been acquired, synthesized, and encoded.
The challenge for the SAnTA research then, was to demonstrate a system that a) had the expertise to select SATs appropriate to analysts' situations, and b) to be useful, that is, analysts, trainers, and facilitators would choose to use SAnTA in their work.

Related Work

Cognitive biases were brought to the attention of the intelligence community by Tversky and Kahneman (1974), among others, in their Science paper, which discussed the heuristics that lead to three biases, representativeness, availability, and adjustment and anchoring. Heuer (1999) notes that intelligence analysts are subject to such cognitive biases, particularly when evaluating evidence, when perceiving cause and effect, and when estimating probabilities. Heuer also notes (page 162) that briefing test subjects on biases and encouraging them to avoid or compensate for bias does not work. He notes (as do Tversky and Kahneman, among others) "Cognitive biases remain compelling even after we become aware of them". Woods (2005) writes that one of the basic support requirements for cognitive work is support for shifting perspectives. This requirement includes *reminding* – suggesting other possibilities as activity progresses, and *critiquing* – pointing out alternatives as activities come to a close. Reminding and critiquing are functions that SAnTA provides analysts as they work with SATs. In addition, SAnTA, as a cognitive tool for analysts, helps analysts explore beyond the SATs of their current focus and to shift and contrast multiple perspectives regarding which SATs may be the most appropriate.
Feltovich et.al. (2004), explore how cognitive engineers can avoid the reductive tendency. In developing SAnTA, we have attempted to avoid the reductive tendency by limiting SAnTA's automation goal. SAnTA's relationship to tradecraft cells (SAT SMEs) is similar to that of ATMs to bank tellers; the tools handle only the simpler, straightforward interactions, while humans must perform the more complex tasks.
Pirolli and Card (2005), point out a number of cognitive biases, and indicate leverage points where these biases might be addressed, based on their sensemaking loop model. Analysts can access SAnTA at each of these leverage points in the analysis process; one of its displays shows which SATs are recommended for each phase of analysis.
In contrast, Hoffman, et.al. (2011) downplay the importance of bias and note the difficulties encountered in applying structured analytic techniques – with or without the tools that are intended to reduce cognitive bias. They propose instead that tools be developed that support perceptual learning and the achievement of expertise. SAnTA may be a small step in this direction. SAnTA enables analysts to explore the range of SATs available to them, it ranks the SATs by their likelihood of applicability based on the analyst's input, and it enables the analyst to justify their selection, based on their situation as articulated in their responses to SAnTA's questions, which they can save for review. The choice of the SAT is the analyst's.
Yet, the intelligence community remains interested in sources of cognitive bias and means of reducing it, as exemplified in IARPA's current ICArUS and Sirius research programs (www.iarpa.gov).

METHOD

We assembled a research team with strengths in analysis, SATs, knowledge capture, and recommender systems. We began our research process by interviewing representatives from the across the analytic community who were aiming to improve analytic tradecraft quality. Using the general interview guide approach, we discussed key needs, approaches, opportunities, and constraints. We confirmed the importance of selecting and applying SATs appropriately and our perception that an SAT recommender could help in addressing the problem. While conducting these interviews we also learned that the likely end users would be not only analysts but also facilitators and instructors.
We decided to conduct applied research, and that our success criteria for the design was that it be useable, useful, and used. We prototyped and evaluated various designs for the user interface. We experimented with alternative representations of domain expertise. We enlisted a second subject matter expert (SME) who is a former analyst and front-line manager of analysts, an instructor of SAT usage, an author or co-author of several popular SAT how-to texts, and, it happens, an articulate advocate for the software. For a programmer we hired a summer intern who turned out to be extremely talented.
We created a series of prototypes (using pure JavaScript for ease of transfer to the end users' security-conscious environments) into which our SMEs input the various factors that influence SAT recommendations, a list of SATs to be recommended, and a matrix characterizing how each factor contributed to the strength of the recommendation of each SAT. The SMEs tested the pre-prototype using historical case studies. We demonstrated this pre-prototype version to the various representatives with whom we had spoken originally, and others, and obtained very positive feedback on our approach. For example:

> *"I have seen many attempts to recommend SATs and they all have failed. This tool works. Let me know how I can help test this in my classes."*

We then partnered with one analytic tradecraft group and supported them as they input their own set of SATs to recommend, their set of factors that influenced their recommendations, and their matrix indicating how each factor contributes to the recommendation of each SAT. The resulting version of SAnTA encapsulates the synthesized expertise of this group in its real-world context. These SMEs are currently checking that SAnTA makes sensible

recommendations by giving SAnTA case studies - previously completed projects - and examining SAnTA's output. This version of SAnTA is also being piloted with analysts to obtain their feedback.

We continue to iterate the software to improve the design of the user interface, the set of factors that influence recommendations, the set of SATs recommended, and the relationships among them. In order to make SAnTA easy to install in the customer's environment initially, we intentionally sacrificed the ability to log user activity - thereby losing the ability to obtain usage analytics as a basis for continuous design improvement. While an instrumented version of SAnTA would be more complex and harder to install and maintain, it would be a significant step forward in the long-term sustainability of the software as a product.

RESULTS

For organizations, SAnTA captures corporate knowledge (synthesized domain expertise) in a visible, accessible, updateable, interactive, and user-focused software tool. For decision-making analysts, SAnTA first requires users to articulate their situation in tradecraft terms, that is, to pause and take a meta-level view of their work in responding to SAnTA's questions. Next, it captures their articulated inputs for later introspection and sharing. Third, SAnTA overcomes limitations of recall, working memory, and availability bias, enabling analysts to go beyond the first SAT that comes to mind and select from the most-likely-suitable ones. Choosing which SAT(s) to use from this set of SAT recommendations prompts further reflection and discovery.

SAnTA inverts the dynamics of selecting a structured analytic technique. In the past, analysts first attended a course or read some material about SATs in a book or online. Then, when the need for an SAT arose, the analyst would recall the SATs mentioned or review a list of SATs, and try to select one that seemed suitable. If all else failed, they could attempt to obtain help from a facilitator. In contrast, with SAnTA, analysts simply describe their situation to SAnTA; characterizing where they are in the analytic process, the resources and constraints in their environment, and so on, and SAnTA provides the user with a small set of recommended SATs, together with a rationale for their selection, the strength of its recommendation, a visualization of where in the analytic process each one applies, and links to further tools or information.

To date, SAnTA research has accomplished several things: first, it documents corporate knowledge and best practices (by requiring their articulation and synthesis), and it does so in a software system that makes this knowledge easily accessible by analysts, trainers, and facilitators. Next, SAnTA contains, and can recommend when appropriate, more SATs than any single facilitator. Third, even if it should turn out to be the case that users need not apply an SAT, in the process of determining this, they will have reviewed a number of questions that every analyst should consider when creating a product. Finally, SAnTA has been developed in a manner that makes it easy to edit or change out the knowledge base it uses.

We have learned first, that when users input data from their past cases, the recommendations are reasonable. Next, ten to twenty questions appear to be sufficient to capture the information needed to make good recommendations (we expect this number could be reduced with further analysis). Third, the current software is easily transferred to and is working well in the analysts' environment. Finally, while the software seems useful as-is, numerous improvements and refinements could be made.

At this point we have buy-in from representatives of the user community. That buy-in has enabled us to obtain funding to transfer the technology to the analytic community and funding to continue the research in a different domain where the approach is also likely to prove fruitful. We have paid attention from the start to creating a product that will be useable, useful, and used, and the feedback we have obtained so far indicates an increasing likelihood of achieving that goal. We have created a representation of domain expertise – for domains with certain characteristics – that makes it relatively easy to capture, synthesize, and make that expertise readily accessible to a broad audience. We have a system that is a shell capable of doing the same for any expertise that can be structured similarly. We have created a system that is expected to help analysts think more thoroughly by making SAT options visible at the time of their application.

DISCUSSION

The challenge for the SAnTA research was to demonstrate a system that a) had the expertise to recommend SATs appropriate to analysts' situations, and b) to be useful, that is, analysts, trainers, and facilitators would choose to use it. Obtaining initial input from the community of users across multiple organizations helped ensure the very first design would be useful, and input to the pre-prototype from two SMEs illustrated the nature of the desired expertise. At this point SAnTA's expertise and utility have been demonstrated to various members of the analytic community and their anecdotal response strongly favors SAnTA's design. When demonstrated to a group of trainers and facilitators SAnTA was seen as useful. That group of domain experts was easily able to express its synthesized expertise in SAnTA's representation. That representation is currently being checked by them using previously completed analytic products.

One of SAnTA's strengths is its user orientation. Instead of thumbing through pages of SATs looking for one whose 'when to use' attributes match their situation, analysts describe their situation to SAnTA and SAnTA generates a set of recommended SATs, together with supporting information that makes it easy for analysts to match the most appropriate technique to their analytic challenge.

Limitations of this report

1. As of this writing, January 2015, actual use by analysts is in progress.
2. As of this writing, January 2015, use by a second group in a second organization is in progress.

Further research on SAnTA

1. Validate the content (factors, weights, and recommendations) using examples from the past.
2. Develop a more-formal means of validating content, especially recommendations.
3. Determine the minimum factors to consider for generating valid recommendations.
4. Find optimal visualizations or ways to present SAnTA's recommendations.
5. Instrument SAnTA and gather usage data to improve the design.
6. Study the usage of SAnTA and the factors influencing the use or non-use of SAnTA.
7. Make it as easy as possible to edit, and to swap in/out, sets of domain knowledge.

Implications

The use of recommender systems for learning (versus consumption) is a neglected area of research. Recommender systems that support informal workplace learning appear to be a potentially valuable tool for the collaborative bootstrapping of expertise. SAnTA demonstrates that recommender systems can support less-experienced naturalistic decision-makers, such as analysts, in specific circumstances.

The conditions of NDM that apply when doing analysis, i.e., making "…tough decisions under difficult conditions such as limited time, uncertainty, high stakes, vague goals, and unstable conditions" (Klein, 2008) make it difficult to pause and think at a reflective or metacognitive level. The presence of an SAT recommender may take pressure off the analyst in that respect, providing a quick solution to the task of selecting suitable, defensible, SATs for use. In these circumstances, giving analysts control, i.e., the ability to change their responses and observe the resulting change in recommendations, together with seeing reasons for the change, enable them to evaluate the specificity, generality, and robustness of the recommendations, to understand the applicability conditions of various SATs, and to begin to internalize this knowledge for later use. The ability of an analyst to save a SAnTA session, including their inputs and the resulting recommendations, enables both users, and those reviewing their work later, to reflect on these perceptions and decisions, and may also help analysts to internalize this knowledge.

To summarize, SAnTA is a knowledge-based recommender system that solicits user input as responses to questions and suggests most-likely-relevant SATs. The user must make the decision of which SAT(s) to use, if any. SAnTA provides brief descriptions of each SAT, explanations of how the user's responses influenced the SAT recommendations, and links to further information on each SAT. SAnTA is written in JavaScript, a design decision that makes it easy to port the software to many environments. SAnTA does have hooks for instrumenting the application and logging how it is used; a future capability that will require SAnTA be linked to a server, which we are currently avoiding for ease of portability. The recommendation algorithm is a simple one. To start with, SMEs rate the applicability of each SAT to each response creating a matrix. As users enter their responses, SAnTA computes the ranking of each SAT based on the current response set. As more responses are entered, SAnTA becomes more certain of its recommendations. This design means SAnTA is easily customized. Sets of SATs and sets of questions and responses are easily exchanged.

REFERENCES

Felfernig, A., Chen, L., & Mandl, M. (2011). RecSys'11 Workshop on Human Decision Making in Recommender Systems. *RecSys'11*, October 23–27, 2011, Chicago, Illinois, USA. ACM 978-1-4503-0683-6/11/10.

Feltovich, P., Hoffman, R., Roesler, A., & Woods, D. (2004). Keeping It Too Simple: How the Reductive Tendency Affects Cognitive Engineering. *IEEE Intelligent Systems*, vol. 19, no. 3.

Heuer, R. (1999). The Psychology of Intelligence Analysis. CIA's Center for the Study of Intelligence.

Hoffman, R., Henderson, S., Moon, B., Moore, D., & Litman, J. (May 2011). Reasoning difficulty in analytical activity. *Theoretical Issues in Ergonomics Science*, Vol. 12, No. 3, 225–240.

Klein, G. Naturalistic Decision Making. *Human Factors*, Vol. 50, No. 3, June 2008, pp. 456–460.

Pirolli, P. & Card, S. (2005). The Sensemaking Process and Leverage Points for Analyst Technology as Identified Through Cognitive Task Analysis. In *Proceedings of International Conference on Intelligence Analysis*. McLean Virginia.

Ricci, F., Rokach, L., Shapira, B., & Kantor, P. (Eds.) (2011). Recommender systems handbook. New York: Springer.

Tversky, A. & Kahneman, D. (1974). Judgment under Uncertainty: Heuristics and Biases, *Science*, Vol. 185, pp. 1124-1131.

Woods, D. (2005). Generic Support Requirements for Cognitive Work: Laws that Govern Cognitive Work in Action. In *Proceedings of the Human Factors and Ergonomics Society 49th Annual Meeting*. Santa Monica, CA.