

The Evolution of User Experiences and Interfaces for Delivering Context-Aware Recommendations to Information Analysts

Adam Fouse^(✉), Ryan S. Mullins, Gabriel Ganberg, and Chad Weiss

Aptima, Inc, 12 Gill Street Suite 1400, Woburn, MA 01801, USA
{afouse, rmullins, ganberg, cweiss}@aptima.com

Abstract. Rigor in the products of information analysis is essential for decision makers to rely on the assessments contained within them. Zelik, Patterson, and Woods [1] defined an eight-attribute metric for communicating the rigor of analytic products. This paper describes two iterations of the process of designing, implementing, and evaluating a context-aware web application that uses this analytic rigor metric to recommend augmentations to analysts' workflow that will improve the quality of the resultant products. We used multiple methods to evaluate this tool with subject matter experts, including brainstorming, collaborative card sorting, semi-structured interviews, cognitive walkthroughs, and heuristic evaluations. This research found that: (1) it is critical to have flexible recommendations that adapt to movements between foraging and sense-making components of workflow, and the changing structure of the analysis; and (2) persistent visualizations of analytic rigor assessments are distracting, and promote interpretation as a performance metric rather than a process aid.

Keywords: Information analysis · Analytic rigor · Human factors · Usability engineering · Systems engineering

1 Introduction

Spurred by the recent and dramatic increase in the volume and variety of available data, the information analysis domain is undergoing something of a renaissance in relation to tools and techniques. This is particularly true of sub-domains that focus on open source information—that information which is publicly available and whose source has little or no expectations that the information will be kept private [2]—including intelligence analysis, financial analysis, legal analysis, and journalism. However, this renaissance has yet to solve the problems germane to the information analysis process. The mistaken assessment that Iraq had weapons of mass destruction [3] is one influential example that has been attributed to failures of process and oversight, among other things.

In this paper, we focus on the problem of rigor in the information analysis process. Zelik and colleagues [1] define rigor as “the appropriateness of [the] fit between analytic processes and contextual requirements”. Their work, discussed in more detail below, provides a framework in which context-aware decision-support systems can be designed to enable and encourage rigorous analytic processes, given an understanding

of the needs of the problem being addressed. Here, we present the findings of two phases of a research and development effort that sought to implement an Adaptive Workspace for Analyst Knowledge and Engagement (AWAKE). AWAKE is a web-based information analysis tool using this framework as a guide. It was designed to support the analysis of open source information by a user who is more akin to an intelligence analyst or journalist than a lawyer or financial analysis, focusing on processing, extracting, and organizing information from primarily textual data. The first phase developed a tool capable of measuring the qualities of a rigorous analytic process given the analyst's interactions with data, and the second phase improved the tool to support a more flexibility in analytic workflows. Usability engineering studies were conducted with subject matter experts throughout each development phase.

2 Background

2.1 Information Analysis Workflows

Argument is the principal tool of information analysts, and, therefore, analytic workflows are designed to enable argumentation. Clark [4] described a systematic approach to information analysis in which the analyst first specifies potential answers to a request and the measures (evidence or grounds) with which they will assess answers, then assesses all of the economical and reliable evidence with a particular eye towards evidence that refutes the one or more potential answers. Clark's method derives from the scientific method, and aligns well with the model of argumentation defined by Toulmin [5], who observed a consistent pattern of reasoning across professions that demand analysis of complex data, professions including business, law, management, and ethics. In this process, expert practitioners bring grounds (evidence) to bear through warrants (theory, itself backed by evidence) to make qualified claims, subject to refutation. Kent [6] proposed that intelligence analysis, one sub-domain of information analysis, was comprised of seven stages: (1) the appearance of the substantive problem; (2) analysis of the substantive problem; (3) the collection of data; (4) the evaluation of data; (5) the moment of hypothesis; (6) additional data collection & hypotheses testing; and (7) presentation. Heuer [7] and later the CIA [8] recommended specific, highly regimented analytic techniques – for example, analysis of competing hypotheses – be inserted into analytic strategies similar to those proposed by Clark and Kent in order to improve the rigor and reliability of the process.

Other analytic theories try to improve on these models by accounting for analysts' intuition and story building skills. Klein [9] asserts that analysts intuit a single, favored claim or action early, rather than systematically enumerate many of them, as Clark prescribes. The analyst then uses mental simulation (a form of story building) and other methods to test the claim and refine the plan of action. Cohen, Freeman, and Thompson [10] described a similar analytic strategy, which they derived from observation of military decision makers and validated in training experiments. They assert that analysts first intuit their favored claim or action given the situation at hand. If time allows, and uncertainty and stakes are high, the analyst builds a story towards the claim, tests it against the evidence, evaluates the story, refines plans, and either cycles back to revise or

replace the story or takes decisive action. Two aspects of this strategy are particularly relevant. First, Cohen and colleagues argue that a story is structured, much like a theory, to address the major aspects of the domain. A story concerning hostile intent, for example, typically addresses the intent of the enemy to do damage, the presence of a target of value to the enemy, the capability of the enemy to locate and attack, and other factors. Second, this strategy recognizes that analysts do and must bend evidence. In any reasonably complex situation, some evidence will refute the analyst's story. The analyst strives to bend (i.e., explain away) that evidence in the testing phase of the process, then assesses whether bending the evidence breaks the story in the evaluation phase.

Most recently, information analysis research has seen a shift towards models that emphasize feedback and iteration, using small adjustments driven by new evidence to hone theories over time. Pirolli and Card's [11] model centers around two loops: (1) a forging loop that includes collecting data, filtering it, and extracting relevant information; and (2) a sense-making loop where analysts actively seek to find the best explanation for those data. During sense-making, analysts establish a frame, which is a hypothesis about how data are connected, and in the presence of new data, they can preserve the frame, elaborate the frame, or establish a new frame if the current frame is no longer plausible. Mellers et al. [12] found that analysts who frequently returned to and incrementally updated their responses to forecasting problems based on changes in evidence were associated with moderate increases in forecast accuracy. AWAKE strives to accommodate flexible, iterative analytic strategies at the macro level, while allowing analysts to use a variety of specific analytic techniques (e.g., ACH, devil's advocacy, red-teaming) within iterations.

2.2 Rigor in Information Analysis Processes

Our research on analytic rigor is grounded in the work of Zelik and colleagues, who define rigor as "the appropriateness of [the] fit between analytic processes and contextual requirements". Their initial research [1] focused on validating their definition of rigor, developing a baseline understanding of an "appropriate fit", and defining a generalized metric for assessing analytic rigor. In this study, each of nine senior analysts was asked observed two junior analysts and determine if their work (i.e., their process and products) was sufficiently rigorous. The study was designed to ensure that one junior analyst would have a high-rigor process and the other would have a low-rigor process, informed by professional intelligence analysts. An elicitation by critique method [13] was used to ensure that responses were non-binary, encouraging the participants to explain why they felt that the process they observed was insufficiently rigorous. The critical finding from this study was that only a small fraction of the senior analyst participants felt that work of even the high-rigor junior analyst was suitable for dissemination to a decision-maker. After analyzing the data from the study, Zelik et al. defined an eight-attribute rigor metric, described in Table 1.

In their follow-on research, Zelik and colleagues [14] sought to define generalizable process measures that feed assessments of analytic rigor, and, using these measures, assess the rigor of analytic processes by tailoring the importance of different attributes to the type of analysis. In this study, participant teams were asked to complete an

Table 1. The eight-attribute analytic rigor metric with definitions, adapted from Zelik, Patterson, and Woods [1].

Attribute	Definition
Hypothesis exploration	The extent to which alternative answers to the analytic problem were considered
Information search	The depth and breadth of the process used to collect relevant data
Information validation	The extent to which data has been cross-referenced and validated
Stance analysis	The extent to which the biases of sources have been acknowledged and placed into a broader context
Sensitivity analysis	The extent to which an analysis hinges on set of possibly unverified assumptions
Specialist collaboration	The extent to which the analyst has incorporated the opinions of other specialists with possibly differing views
Information synthesis	The extent to which the analyst has gone beyond simply stating the data they reviewed and toward making inferences that explain their answer to the problem
Explanation critiquing	The extent to which other perspectives were incorporated and presented in the final product

analysis under one or two experimental conditions that placed the participants under different types of stress, e.g., communicating in a collocated or distributed environment, or under the time pressure of an impending deadline. Verbal transcripts and the analytic products of these studies were manually coded by two individuals. The coders were asked to rate each team as either low, moderate, or high in each applicable attribute, and provide a rationale for their rating. An analysis of the coders' rationale provided insight into the specific measures that should be used for assessing analytic rigor. For example, when rating the information search attribute, the two most common measures were the number of documents reviewed (i.e., less for novice analysts than for expert analysts) and the time spent reviewing those documents (i.e., expert analysts spent more time reviewing the documents than novices). AWAKE builds on this work by defining a set of semantic process measures (e.g., the number of hypotheses associated with an analytic problem), and implementing an automated rating system that aggregates individual measures into attribute-level scores. These ratings and scores are used to derive and deliver recommended changes to the analyst's process.

3 The AWAKE Application

AWAKE is a web-based application for analysis of textual open source information. Target users include intelligence analysts (e.g., analysts at the Central Intelligence Agency), citizen investigative journalists (e.g., members of Bellingcat¹), business

¹ <https://www.bellingcat.com>.

analysts, or legal analysts, though the potential user base could be much larger. The design of the application has two high-level goals: (1) to recommend changes to analytic process that will make the resulting products more rigorous; and (2) to enable a flexible and collaborative information analysis process, where teams of analysts contribute asynchronously to the development of analytic products. The specific design and development process related to each of these goals is discussed below.

To achieve these goals, AWAKE must have contextual awareness that enables an understanding of the current state of an analysis, in order to assess analytic rigor. To this end, AWAKE defines a context model [15], shown in Fig. 1, that provides data types, stored as entities in the graph, for analytic problems, hypotheses (i.e., the set of possible specific answers to the analytic problem), areas of analysis (i.e., the general categories of search and analysis that may lead to the development of hypotheses), documents, and snippets (i.e., the relevant information extracted from documents). The context model defines explicit, named, and semantically meaningful relationships between these entities, for example, hypotheses are connected to evidence through a “HasEvidence” relationship. AWAKE’s analytic rigor assessments, discussed below, use the current state of the context model, and the history of interactions that led to that state, to determine the attribute scores the drive recommendations.

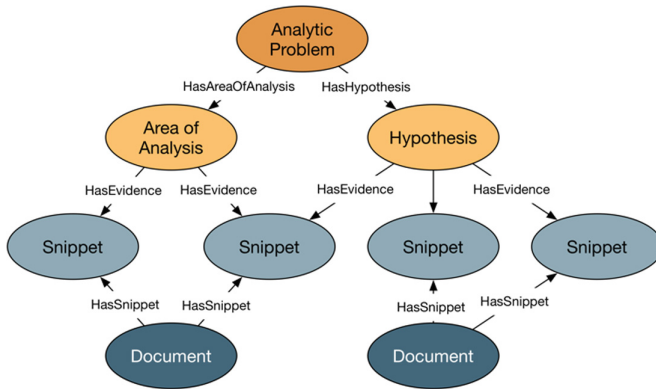


Fig. 1. AWAKE’s context model. Ovals with *orange* hues refer to entities representing constructs of the analytic process, and ovals with *blue* hues refer to entities representing the information leveraged in the analytic process. Relationships connecting the entities are shown with their semantic labels.

AWAKE is built on the Context Platform, a suite of tools developed to enable the development of context-aware applications. At the core of the Context Platform is a graph-based knowledge store, which implements AWAKE’s context model and is the primary data persistence layer. The knowledge graph is supported by other data stores that specialize in supporting full-text search and spatiotemporal search. The current implementation of the Context Platform uses the Titan graph database and Apache Cassandra for the knowledge graph, supported by Elasticsearch for full-text and spatiotemporal search. The Context Platform provides a Gremlin-based scripting interface

to analyze the data in the knowledge graph. Analytics can be set to run continuously, at specific time intervals, or on-demand. Data is integrated into the Context Platform’s data store by either streaming or batch-uploading data to its RESTful web services or command line interface. The AWAKE web application interfaces with the Context Platform using a JavaScript client library that provides a WebSocket connection to the server.

4 Phase 1: Designing to Assess Rigor

The first version of AWAKE (Fig. 2) was designed primarily to address our goal of recommending changes to processes that improve analytic rigor. With the assumption that the tool would support the analysis of textual open source information with an initial focus on the sub-domain of intelligence analysis, we conducted several in-person and distributed brainstorming sessions with former intelligence analysts to define the requisite baseline functionality. These requirements included the ability to: (1) define an analytic problem; (2) define hypotheses and areas of analysis to anchor and guide the collection and analysis of information; (3) find documents using search-engine style querying; (4) explore a document and extract critical information from the text of documents; (5) rate documents and extracted information based on their agreement, as well as standard scales for reliability and credibility [2]; (6) link extracted information to hypotheses and areas of analysis as evidence; (7) develop written assessments for hypotheses, areas of analysis, and analytic problems; and (8) rate their subjective confidence in an assessment. The interface was designed based on a three-pane structure. The left-most pane provided contextual information about the analytic task. The center pane showed the content that was being focused on in the current analytic task. The right pane provided tools and relevant information that may make that task easier. Looking at the example shown in Fig. 2, the analyst is exploring a document related to an analytic problem. On the left is the analytic problem (at the top), a visualization of the analytic rigor assessment immediately below, which are persistently shown, and a list of the current hypotheses related to that problem. In the center is the document itself, with highlighted text indicating extracted information that can be linked to hypotheses as evidence. On the right is an interface where the user can add a tag to the document, and a list of all the extracted content from the document. AWAKE provides a suite of standard mouse and keyboard interactions for manipulating data, including the use of button, drop-down lists, toggles, text input areas, and drag-and-drop features to link evidence.

User interactions are the primary input to the analytic rigor measurements. For example, one measure of rigor related to the information search attribute observed by Zelik et al. [14] was the amount of time spent reviewing documents, which can be measured as the difference between the time they open and close a document in AWAKE. Building on this, we developed similar measures, given in Table 2, that were generated through brainstorming exercises, and categorized using a collaborative card sorting exercise. We employ different strategies to calculate these measures in the AWAKE system, including heuristics (e.g., a minimum number of hypotheses to be considered given an analytic problem), averages (e.g., the average credibility score),

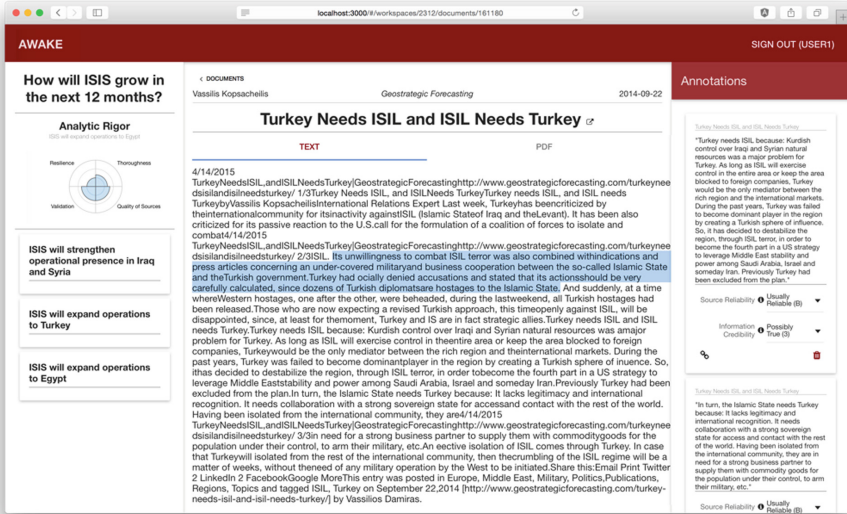


Fig. 2. The first iteration of the AWAKE user interface, showing the in-text highlighting capability (*text with blue background*) that enables information extract from a document. The interface also displays contextual information about the analysis (*left sidebar*), and the tags and extracted information (i.e., annotations) associated with the document (*right sidebar*).

and others. Given the limited capabilities of the AWAKE prototype, we chose to only implement measures related to the hypothesis exploration, information search, information validation, and sensitivity analysis attributes of rigor.

The final piece to the first iteration of the AWAKE puzzle is the delivery of process recommendations. We designed these recommendations to address deficiencies in the individual measures defined above. Recommendations were also design to be targeted, providing a specific action that could be performed by the user on a specific piece of information. These were formulated as generalized critical thinking prompts, which Cohen et al. [10] found to increase the number of hypotheses generates, increase the amount of conflicting evidence identified, induce resolution of conflicting evidence, and improve the accuracy of resultant decisions. For example, when the average credibility of the evidence associated with a hypothesis is low, AWAKE can deliver two types of prompts to the user: (1) “Can <low-credibility evidence> be corroborated by any of this evidence that is associated with this hypothesis?”; and (2) “Can you provide a credibility rating for <unknown-credibility evidence>?”. Recommendations are accessed by clicking on the sectors of the Analytic Rigor graph (shown in the upper-left corner of Fig. 2), which presents the recommendations as a list, along with details about current rigor score for that attribute, given a target element of the analysis (i.e., an analytic problem or hypothesis).

Table 2. Example measures of analytic rigor implemented in the AWAKE prototype system. Measures are categorized by rigor attribute, and associated with a measurement strategy.

Attribute	Measure	Example recommendation
Hypothesis Exploration	Count of hypotheses	“Should you another hypothesis?”
	Count of evidence linked to a hypothesis	“Can you find information relevant to this hypothesis in any of these unread documents?”
Information Search	Time in search UI	“Have you thoroughly reviewed all relevant information?”
	Average source reliability score	“Does this < source>’s information make them more reliable?”
Information Validation	Average information credibility score	“Can you rate the credibility of this information?”
	Agreement between sources	“Do these two pieces of information support or refute each other?”
Sensitivity Analysis	Analysis of robustness of credibility, reliability, and agreement	“Are you relying too heavily on this information?”

4.1 Evaluation and Results

The first iteration of AWAKE was evaluated as a longitudinal study that combined semi-structured interviews [16], heuristic analysis [17], and cognitive walkthroughs [18]. The study began with a focus-group style semi-structured interview where three former intelligence analysts, who would serve as the participants throughout the study, reviewed the functional requirements for the AWAKE system, the measures and measurement strategies for assessing analytic rigor. Their verbal feedback was recorded and used to drive the design and implementation of the AWAKE user interface mock-ups.

Approximately four months later – once the baseline analysis definition, search, document exploitation, and rigor analysis capabilities were implemented – these same participants were given accounts on the system and asked to respond independently to an analytic problem concerning aircraft technology. They were provided with a baseline dataset to draw from, but otherwise were free to conduct their analysis using the workflow that fit best for them. This phase of the evaluation lasted for approximately eight months. During this period, we conducted four focus group sessions where we gathered bug reports and feature requests, and conducted heuristic assessments of the tool using Gerhardt-Powals’ [17] framework.

This evaluation period ended with a cognitive walkthrough [18] of the AWAKE tool. The walkthrough had participants walk the facilitator through their current analysis. The majority of the time was spent walking through and discussing three areas – information extraction, evidence linking, and sense-making of the analytic rigor scores – that had been the cause of friction during the evaluation process. The walkthrough ended with a final heuristic assessment.

Two major findings from this evaluation influenced the second phase of the research and development effort. First, that the persistent visualization of the analytic rigor assessment was not useful. It was visually distracting, and analysts spent more time trying to figure out what their score was and why it changed than acting on the recommendations hidden behind it. From a process perspective, the analysts felt that the visualization, with its emphasis on scoring, was interpreted as a performance metric, rather than a process aid, which would cause anxiety in the workforce. Second, that the interface layout was too restrictive. The first version of the interface was designed around the idea of foraging and sense-making loops [11], but interfaces to support the two activity loops were separated and tedious to move between. This didn't reflect the dynamic nature of analysis, and limited the flexibility with which participants could organize their analysis. For example, to move a group of documents from an area of analysis to a hypothesis required several steps and an excessive number of mouse clicks. The participants also experienced confusion as to their place in the analysis due to design similarities for interfaces representing different concepts (e.g., the use of the same card style to represent analytic problems, hypotheses, and areas of analysis).

5 Phase 2: Designing to Enable Flexible Processes

Based on the findings from the first evaluation, we began work on the second phase of AWAKE development. A critical difference in this phase of development was the existence of a validated context model and analytic rigor assessment capability. As such, the majority of the effort in this phase was spent redesigning the interface to provide a more fluid user experience. The rigid three-pane structure of the previous version was changed to a two-pane model, with a goal of enabling more fluid transitions between foraging and sense-making by using a more dynamic structure. The updated design unified interfaces for search results and for user-defined collections of documents associated with areas of analysis and hypotheses. This enabled viewing documents and structure side-by-side, which addressed both the need to fluidly move between analysis loops and the need to easily re-organize analyses.

In this updated design, shown in Fig. 3, information presented on the left provides background information and context for the information shown on the right. Similarly, the information shown in the bottom half of either side of the interface provides context for the information shown on the top. The analytic problem currently addressed by the analysts is always shown as a card in the top-left of the screen, and can be expanded to show the associated hypotheses, areas of analysis, and analytic rigor recommendation. Rigor recommendations are now shown in a simple list with attribute labels serving as dividers. Many features from the first design were reworked to fit into this new interface, including the search interface, search results list, and document viewing and exploitation features.

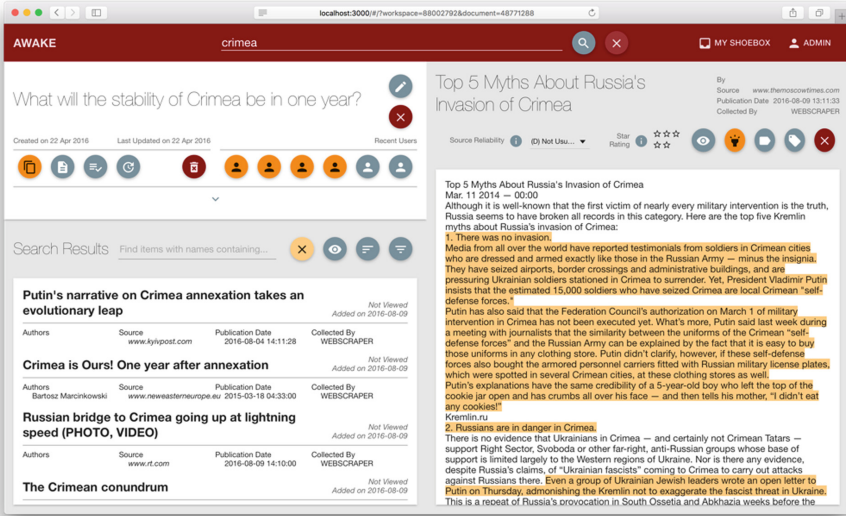


Fig. 3. The second iteration of the AWAKE application, showing the current analytic problem (upper-left), search results (bottom-left), and an exploited document (right).

5.1 Evaluation and Results

The second iteration of the AWAKE application was evaluated using only a cognitive walkthrough. Two of the participants from the first study were walked through the new user interface design, and asked to conduct short analysis that addresses the analytic problem “What will the stability of Crimea be in one year?”. They were given a static set of possible hypotheses – more stable, less stable, or no change – and asked to find and exploit information related to each of the hypotheses. Their interactions with the system were captured, and we written records of the follow-on discussion were collected.

The results of this evaluation were generally positive. The analysts found that the restructuring of the analytic rigor recommendations alleviated their concerns regarding its use as a job performance metric, and the new list format provides equivalent utility without the visual distraction. They also indicated an improvement in the workflow, specifically because of the reduced confusion about what information they were looking at. However, they mentioned that the left-right and top-bottom arrangement of the interfaces took considerable adjustment. For example, when the analyst has selected an analytic problem and executed a search query, the analytic problem is shown at full-height on the left and the search results are shown at full-height on the right. When they select a document from the search results, the left side of the screen instantly splits, with the top half showing the analytic problem and the bottom half showing the search results, and the selected document being shown at full-height on the right. Further discussion revealed that their confusion was likely due to the inability to predict or track how the information was moving on the screen. One solution that is being explored is to animate the content as it resizes and moves into position.

6 Discussion and Future Work

As described above, the AWAKE application has shown potential as an information analysis tool. We have completed two formative usability evaluations that have provided, with a representative, if small, population of potential users. Future work on the AWAKE application will focus on evaluating the application as part of human-in-the-loop performance evaluations, compared to industry-standard analytic tools (e.g., Analyst's Notebook and nSpace 2). Additional work will explore more robust analytic rigor measures that leverage machine learning tools to automatically tailor the definition of "appropriate" based on the characteristics of the analytic problem, and evaluate the performance differences between analysts using AWAKE with and without the support of the analytic rigor recommendations.

Acknowledgments. The research reported in this document was performed in connection with contract FA8750-14-C-0124 with the U.S. Air Force Research Laboratory. We would like to acknowledge Dr. Caroline Ziemkiewicz, Ms. Stacy Lovell Pfautz, Dr. David Woods, Ms. Valerie Champagne, Ms. Amy Rider, Mr. David Johnson, and Mr. Roger Dziegiel for their contributions to this research as thought partners and subject matter experts.

References

1. Zelik, D.J., Patterson, E.S., Woods, D.D.: Understanding rigor in information analysis. In: 8th International Conference on Naturalistic Decision Making, Pacific Grove, CA (2007)
2. U.S. Department of the Army: Open-Source Intelligence. Army Techniques Publication 2-22.9. U.S. Department of the Army, Washington, DC (2012)
3. Robb, C.S., Silberman, L.H., Levin, R.C., McCain, J., Rowen, H.S., Slocombe, W.B., Cutler, L.: The Commission on the Intelligence Capabilities of the United States Regarding Weapons of Mass Destruction: Report to the President of the United States. Commission on Intelligence Capabilities Regarding WMD, Washington, DC (2005)
4. Clark, R.M.: Intelligence Analysis. American Literary Press Inc, Alexandria (1996)
5. Toulmin, S.: The Philosophy of Science: An Introduction. The Hutchinson University Library, London (1953)
6. Kent, S.: Strategic Intelligence for American World Policy. Princeton University Press, Princeton (1965)
7. Heuer, R.J.: Psychology of Intelligence Analysis (1999)
8. United States Government: A Tradecraft Primer: Structured Analytic Techniques for Improving Intelligence Analysis, Washington, DC (2009)
9. Klein, G.: Intuition at Work: Why Developing Your Gut Instincts Will Make You Better at What You Do. Currency Books, New York (2002)
10. Cohen, M.S., Freeman, J.T., Thompson, B.T.: Critical thinking skills in tactical decision making: a model and a training method. In: Canon-Bowers, J., Salas, E. (eds.) Decision-Making Under Stress: Implications for Training and Simulation. American Psychological Association Publications, Washington, DC (1998)
11. Pirolli, P., Card, S.: The sensemaking process and leverage points for analyst technology as identified through cognitive task analysis. In: Proceedings of the International Conference on Intelligence Analysis (2005)

12. Mellers, B., Stone, E., Murray, T., Minster, A., Rohrbaugh, N., Bishop, M.: Ungar, L.: Identifying and cultivating superforecasters as a method of improving probabilistic predictions. *Perspect. Psychol. Sci.* **10**, 267–281 (2015)
13. Miller, J.E., Patterson, E.S., Woods, D.D.: Elicitation by critiquing as a cognitive task analysis methodology. *Cogn. Technol. Work* **8**, 90–102 (2006)
14. Zelik D.J., Patterson E.S., Woods D.D.: Measuring attributes of rigor in information analysis. In: *Macroognition Metrics and Scenarios: Design and Evaluation for Real-world Teams* (2010)
15. Pfautz, S.L., Ganberg, G., Fouse, A., Schurr, N.: A general context-aware framework for improved human-system interactions. *AI Mag.* **36**, 42–49 (2015)
16. Cooke, N.J.: Varieties of knowledge elicitation techniques. *Int. J. Hum. Comput. Stud.* **41**, 801–849 (1994)
17. Gerhardt-Powals, J.: Cognitive engineering principles for enhancing human-computer performance. *Int. J. Hum. Comput. Interact.* **8**, 189–211 (1996)
18. Wharton, C., Bradford, J., Jeffries, R., Franzke, M.: Applying cognitive walkthroughs to more complex user interfaces: experiences, issues, and recommendations. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM (1992)

Advances in Usability and User Experience
Proceedings of the AHFE 2017 International
Conference on Usability and User Experience, July
17-21, 2017, The Westin Bonaventure Hotel, Los
Angeles, California, USA
Ahram, T.Z.; Falcão, C. (Eds.)
2018, XV, 711 p. 278 illus., Softcover
ISBN: 978-3-319-60491-6