RC24140 (W0612-065) December 15, 2006
Computer Science

IBM Research Report

A Study of Information Gathering and Result Processing in Intelligence Analysis

David Gotz, Michelle X. Zhou, Zhen Wen
IBM Research Division
Thomas J. Watson Research Center
P.O. Box 704
Yorktown Heights, NY 10598

Research Division
Almaden - Austin - Beijing - Haifa - India - T. J. Watson - Tokyo - Zurich

A Study of Information Gathering and Result Processing in Intelligence Analysis

David Gotz
dgotz@us.ibm.com

Michelle X. Zhou
mzhou@us.ibm.com

Zhen Wen
zwen@us.ibm.com

IBM T. J. Watson Research Center
19 Skyline Drive
Hawthorne, NY 10532 USA

ABSTRACT

The development of new interface tools for intelligence analysis requires a thorough knowledge of the analytical process, and a comprehensive understanding of analyst activities that can best benefit from intelligent user interface technology. However, due to the secretive nature of the domain, there have been relatively few public studies examining analyst behavior. In this paper, we present the results from a user study designed to explore the ways in which analysts gather and process information. The study asked three subjects to perform two different analytical tasks. The tasks were designed to gather observations from two points in the analytical process: (1) the beginning of an investigation, where little context is available, and (2) the middle stages of the analytical process, where there can be significant context based on gathered evidence and developed target models. We discuss both the study methodology and key results. Our results include observations of how analysts formulate information requests and process the results.

1. INTRODUCTION

The analytical process for Intelligence Analysis (IA) is both difficult and complex. The enormous volumes of data and the complexity of the logical tasks to be performed contribute to making IA a challenging problem to which intelligent user interaction (IUI) technologies can be brought to bear.

However, before technological solutions can be designed and deployed, it is important to first understand the unique demands and requirements inherent in the analytical process as applied to intelligence gathering and understanding. Significant advances will be difficult to achieve without first building a solid foundational understanding of the challenges faced by analysts and the limitations of their existing tool-set.

As part of this effort to understand how IUI tools can best be applied to the IA domain, we have completed a user study designed to provide insights into the IA process itself and to explore the ways in which analysts gather and process information over the course of their investigation. Our study takes a cognitive task analysis approach to the problem. We prepared two distinct analytical tasks to observe both the beginning and middle stages of the IA process. The first task, at the start of the IA process, allowed us to observe how analysts formulate and issue information requests when there is relatively little available context. The second task, at the middle of the IA process, provided us with a view of analytical behavior in the presence of a well-defined context including both gathered evidence and a structured target model.

In this paper, we discuss both our study methodology and key results. We begin with a review of related work in Section 2. Section 3 describes our study design and goals. We discuss the results from our study and highlight a variety of interesting responses in Section 4. Finally, in Section 5, we conclude with several observations regarding IA and how IUI technologies can best be leveraged.

2. RELATED WORK

Historically, there have been relatively few public studies that have examined the role to which advanced interface technology can aid analysts in there tasks. However, new technologies that can generate more vast data collections, coupled with recent geopolitical events, have led to a public recognition of the need to explore how user interface technology can be applied to the IA domain.

There are several valuable resources available the examine the analytical process in general. These include a number of books published former government intelligence analysts [4, 5, 7]. These books provide valuable insights into the complex analytical process as seen by those who practice it.

However, only a small number of scientists external to the intelligence agencies have published peer-reviewed studies aimed at gaining a better understanding of the core technological challenges faced by IA tools. One of the early studies [9] explored the analytical process using cognitive task analysis and developed a series of design concepts aimed at combating the problem of data overload.

These concepts include two key areas where IUI technology may be exploited: (1) identification of high-profit documents, and (2) recognition of evidence that may alter previous hypotheses This work focuses on the foraging loop where analysts are sifting through large amounts of information. In contrast, our study concentrates on how the analyst, given a current context, organizes their analysis and formulates subsequent information requests.

The Novel Intelligence from Massive Data (NIMD) program has sponsored both observational studies and captured trace studies from instrumented analysis environments [11]. The electronic traces are captured via Glass Box Analysis [3], a technique that instruments existing computer-based analytical tools to capture traces of analyst behavior during live IA activities.

The NIMD studies indicate that analysis is tightly intermingled with other tasks, such as data collection and report generation. This

is in contrast to the traditional model of IA as several discrete steps. These studies have led to proposed metrics for the evaluation of intelligent technologies [8].

In other research, scientists have explored alternative IA models and performed user studies that suggest how technologies can be inserted into these models. Pirolli and Card [10] examine a framework in which IA is broken into two loops: a foraging loop and a sense-making loop. They then enumerate several leverage points within those loops where technology could be used to assist in the process. Badalamente and Greitzer [1] have enumerated a "top ten" list for analytical tool development after running a one-day workshop in which they worked with analysts in both the counter-intelligence and counter-terrorism disciplines.

A key difference between our work and these two studies [1, 10] is our focus on the role of an explicit representation of context in the analytical process. We show that structured target models [2], either developed ad-hoc by an analyst or provided as part of the task, are embraced by analysts to focus and organize their investigation. These models, representative of the current context of an analysis, can then be used as input to IUI algorithms.

3. USER STUDY DESIGN

We conducted a three-hour user study with a group of three participants. Over the course of the study, we presented each participant with two distinct analytical tasks. We observed their behavior and solicited explicit feedback through both verbal interviews and written questionnaires. In this section, we describe the subject pool, task specifications, and observation mechanisms in more detail.

3.1 Study Participants

The user population in our study consisted of three participants with a range of analytical experience. Two of the three subjects have extensive experience as government intelligence analysts. They have worked in various agencies for varying lengths of time and therefore posses a breadth of practical experience in IA activities. The third participant, an independent researcher from within our lab, has no experience in intelligence analysis but significant experience in tackling difficult research problems.

3.2 Task Specifications

Recognizing that the three hour duration of our user study would be inadequate for completing a complex investigation, we designed a two-stage study with a pair of smaller exercises corresponding to specific stages in the analytical process. Users were given one hour to perform each of the two exercises. After each exercise, we gathered feedback via interviews and questionnaires as described in Section 3.3.

3.2.1 Task One: Initial Investigation

The first task in our study was designed to observe analyst behavior during the initial investigation stage of an analysis. Analysts were provided with excerpts from a letter written by Enron whistleblower Sherron Watkins. The letter described a number accounting practices which she considered suspicious.

With this as their initial information, the analysts were asked to uncover as much as possible about any accounting irregularities at the company. Analysts were asked to work individually, with no communication allowed between participants. Analysts were given access to three additional sources of information:

- **E-Mail:** A keyword-based search engine providing access to roughly 500,000 actual Enron emails released to the public by the Federal Energy Regulatory Commission [6].

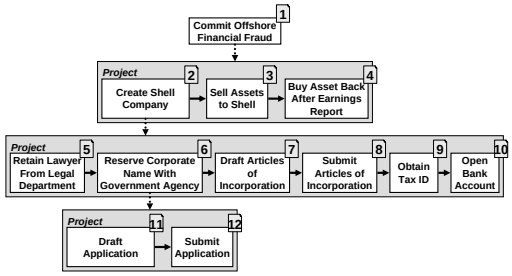

Figure 1: The process model provided to analysts for task two.

- **Chat Transcripts:** A keyword-based search engine providing access to computer generated chat transcripts with "Enron-like" topic and term distributions. These chats were artificially generated using the email corpus as seed material.
- **The Internet:** Public search engines and web sites, allowing access to a wealth of publicly available information on Enron and the surrounding investigations.

The analysts were asked to begin their investigation and given one hour to make progress. The analysts were asked to record notes related to their investigation and their behavior over the course of the task was carefully observed.

3.2.2 Task Two: Hypothesis Modification

The second task in our study was designed to observe the way in which analysts work in more advanced stages of an investigation, after a set of hypotheses have already be developed. In developing this task, we aimed to observe the analytical behavior associated with hypothesis creation, verification, and modification.

In this task, Analysts were provided with invented information from a fictional informant. The informant claimed to have inside information regarding suspected fraudulent activity. Furthermore, the analysts were told that the investigation was already underway. Rather than starting from scratch, the analysts were taking over for a previous investigator and were provided with well-developed model-based hypotheses regarding the suspected activity. The hypotheses were provided in the form of a system model (see Figure 1) as described by Robert Clark in his book on target-centric intelligence analysis [2].

In addition, the analysts were provided with specific pieces of evidence that had already been identified. The analysts were tasked with continuing the ongoing investigation, including verification of existing hypotheses, discovery of new information, modification of old hypotheses, and development of new hypotheses. Unlike the first task, the participants were allowed to work collaboratively.

The analysts had access to the same three sources of information as in task one: the Enron email corpus, generated chat transcripts, and general Internet access. Using the provided process model and associated evidence, the participants were asked to continue the investigation. The analysts were asked to record nodes related to their investigation and the team of organizers carefully observed their behavior.

3.3 Gathering Feedback

We gathered feedback from the analysts through several mechanisms. Two methods have already been discussed: written notes by the analysts, and behavior observations. In addition, two other tools were used to gather more specific feedback from the analysts: questionnaires and interviews.

Each analyst was asked to answer three questionnaires during the user study. First, each analyst completed a pre-study profil-

ing questionnaire to gather background information. Task-oriented questionnaires were then administered at the end of each phase.

We also performed one-on-one interviews with each analyst to more deeply probe their activities. Initial questions were taken directly from the post-task questionnaire and follow up questions were developed based on the analysts' performance and opinions.

4. STUDY RESULTS

In this section, we report the key results from our study. Our analysis highlights a number of important factors surrounding analyst behavior in both information gathering and result processing.

4.1 User Profiles

Each participant in our study completed a pre-task questionnaire. Conclusions based on the self-reported results are presented below.

Investigative Style: Participants were asked to self-identify their approach to IA as either top-down (formulating a full model before collecting data) or bottom-up (collecting as much information as possible before developing a model). All participants expressed that they employ a mixture of techniques. One respondent declared they did not fit into any category, claiming to use "whatever makes sense to understand and solve the problem." The heterogeneity in investigative techniques matches the findings of several other studies. In particular, Pirolli and Card [10] report their observations show techniques are applied in an "opportunistic mix."

These findings indicate that both information gathering and result processing are highly integrated and occur "on-demand" within the IA process. This is an important design point for the development of IUI tools, which must be designed to avoid overly constraining the analytical process so that analysts are not forced to conform to any single investigative model.

Record Keeping: Asked about their record-keeping practices during an investigation, all respondents mentioned unstructured notes. Intriguingly, none of the experienced analysts mentioned computer aided mechanisms for organizing their investigation. Only the research scientist mentioned note taking on a computer and capturing "links" to information. The reliance on unstructured notes again indicates the organic process by which analysts merge information gathering tasks and result processing. However, the results in Section 4.3 indicate that there is significantly more structure in the process than the analysts admit in their pre-task questionnaire.

Technology: We also asked the participants for their opinion regarding how technology can aid the IA process. Each participant was provided with a list of options and asked to select and rank which tools would be most useful. Given highest priority were computer tools designed track and generate hypotheses throughout the life of an investigation. This finding is consistent with other studies which placed such a tool high on their own list of the top ten needs for IA tool development [1]. Also ranking highly in both our study and [1] is the need for improved collaboration for joint analysis.

4.2 Inhibiting Factors

After each task in our study, we asked the analysts to identify the primary factors that inhibited their investigation. We identified several common themes based on their feedback.

Additional Tools: Information storage and browsing tools were identified as particularly important due the electronic nature of the data sources. Because information dismissed as unrelated often proved to be more important later in their investigation, the analysts wanted tools that enabled context-sensitive queries over the history of retrieved information. Such IUI tools could allow analysts to easily re-discover previous work.

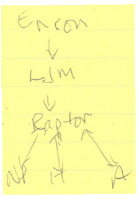

Figure 2: An analyst's hand-drawn entity model.

Another need identified by the analysts is a tool for piecing together electronic information. The process of moving between electronic sources and a paper investigation record proved onerous, indicating a strong need for intelligent applications that can help organize information within the context of an investigation.

Information Sources: All participants talked about the need for rich and reliable information sources. Framed broadly, there were two categories of information sources: (1) high level background knowledge and (2) low-level event or factual information. Both types of resources are essential tools for completing an analysis.

In task one, corresponding to the early stages of an investigation, the analysts relied on high-level information to develop domain-specific models. In task two, the users spent more time looking for low-level information as they gathered evidence to support or refute the already-developed process model.

Our observations indicate that, despite the shift in focus between tasks, analysts tightly couple their access to both data source types. As the analysts discovered potential evidence, they would immediately access higher-level information sources to improve their understanding of what they had found. This hints that IUI tools for identifying relevant data sources and integrating them into a uniform environment would be extremely useful.

Domain Expertise: Another theme voiced by the analysts in our study is that of domain expertise. Understanding the terminology associated with the Enron investigation was a difficult task and all participants discussed the desire for direct collaboration with domain experts. This highlights the collaborative nature of the information gathering process and indicates a need for intelligent tools capable of connecting analysts to relevant experts.

The creating of analysis templates, built by domain experts, also had broad interest from the analysts in our study. These templates, designed for specific domain problems, could be used to boot-strap an investigation in the early stages to provide domain novices with formal models and techniques developed by experts.

4.3 Benefit of a Structured Model

Structured target models were important parts of both tasks in our study. In the first study, we observed that the analysts developed their own ad hoc models. In the second study, a mature process model was provided as part of the analyst's initial information.

4.3.1 Ad Hoc Model Development

In the first task, the analyst's initial information was provided in the form of an unstructured text document. We observed that the study participants began to develop structured models using pencil and paper. The model allowed them to organize terms and concepts into a semantic web as relationships were discovered. For example, Figure 2 shows the entity model diagrammed by one of the analysts. In the follow-up interview, the analyst stated that the model allows her to create a "big picture view" of the problem from which she could begin a more detailed investigation for each item.

The ad hoc model served as the organizational entity around which the IA process was performed. Both the model's structure

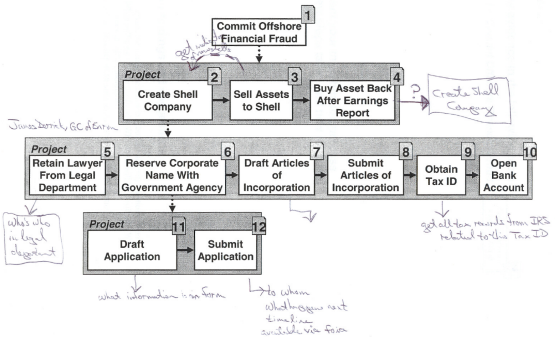

Figure 3: Analyst annotations to the provided process model.

and its contents were used to develop subsequent requests for information. The analysis of newly acquired results was expressed in part by updating the model to reflect new conclusions.

4.3.2 Process Models

A more formal model was introduced in the second task. Each participant was provided with a process model description of the hypothesized activity. The model served as the launching point for many of the information gathering actions performed by the analysts.

For example, analysts were seen annotating the process model with identified evidence, important queries, and key questions that needed to be addressed. During the post-task interviews, the process model was said to aid in laying out the larger context and focusing the investigation on the most important elements.

It is important to note that the analysts used the process model for more than just organizing evidence and queries. More significantly, they began to alter the model by adding new elements and relationships. The changes made by one participant are shown in Figure 3. Our observations indicate that analysts are open to formalizing their analytical thought process by explicitly altering the model as they draw conclusions.

We believe that an effective computer tool that enables the intuitive construction of these models can be very powerful in the hands of an analyst for the organization of both gathering and processing results. In addition, given that these models correspond to the current state of the analysis, they provide a representation of the context which can be exploited by IUI algorithms designed to aid in the investigation.

4.4 Collaboration

Our study encouraged collaboration between the participants during the second task. As indicated in Section 4.2, collaboration can aid information gathering and processing by connecting analysts with domain experts. In addition to providing basic background knowledge, collaboration enables the sharing of both collected evidence and the models developed as part of result processing.

However, there are disadvantages associated with collaboration as well. In particular, the danger of *groupthink* is far greater when people are working together as a larger unit. Groupthink is the notion that groups can make irrational decisions because individual members feel the need to conform.

For example, during the second task of our study, where collaboration was encouraged, we observed less variety in the information gathering and result analysis process. This danger highlights an additional role where IUI tools have great potential: the detection and user notification of alternative hypotheses that have not been given adequate attention in the analysis.

5. CONCLUSIONS

We have conducted a user study designed to explore the ways in which analysts gather and process information. The study asked a group of three participants to attempt two distinct analytical tasks. The first task asked the analysts to begin an new investigation, allowing us to observe their behavior during the initial stages of the analytical process. The second task asked the participants to take over an investigation which was already well underway, allowing us to examine how analysts work with existing information and formally modeled hypotheses.

Based on observations, interviews, and questionnaires, our investigation has identified a number of important aspects of the information gathering and result synthesis process. The results from our study can be used to help build the foundation of knowledge surrounding the intelligence analysis process that is an essential prerequisite to building effective intelligent user interface tools.

Among the most critical observations, our results show that information gathering and result processing are tightly coupled into a single iterative stage. In addition, the study indicates a similar relationship between high-level and low-level information sources as analysts quickly move between gathering background knowledge and evidence seeking.

Perhaps most importantly, our study indicates that the development of formal models are often an inherent part in the analytical process. By providing tools that make the explicit building of these models easy and efficient, we can create technologies that both directly benefit the analyst and formally capture the context of the investigation.

Taken together, the observations presented in this paper help build a better understanding of how analysts perform their work during the intelligence gathering and result processing stages of an IA investigation. This foundation will help us understand how IUI tools can be developed and where they can be best applied to aid in the IA process.

6. REFERENCES

[1] R. V. Badalamente and F. L. Greitzer. Top ten needs for intelligence analysis tool development. In *Proceedings of the International Conference on Intelligence Analysis*, May 2005.

[2] R. M. Clark. *Intelligence Analysis: A Target-Centric Approach*. CQ Press, July 2003.

[3] E. Hampson and P. Cowley. Instrumenting the intelligence analysis process. In *Proceedings of the International Conference on Intelligence Analysis*, 2005.

[4] R. J. Heuer, Jr. *Psychology of Intelligence Analysis*. U.S. Government Printing Office, 1999.

[5] M. D. Jones. *The Thinker's Toolkit: 14 Powerful Techniques for Problem Solving*. Three Rivers Press, 1998.

[6] B. Klimt and Y. Yang. Introducing the enron corpus. In *Proc. of the Conf. on Email and Anti-Spam*, Mountain View, CA, July 2004.

[7] L. Krizan. *Intelligence Essentials for Everyone*. Diane Pub Co, 1999.

[8] E. Morse, M. P. Steves, and J. Scholtz. Metrics and methodologies for evaluating technologies for intelligence analysts. In *Proceedings of the International Conference on Intelligence Analysis*, 2005.

[9] E. S. Patterson, D. D. Woods, D. Tinapple, and E. M. Roth. Using cognitive task analysis (CTA) to seed design concepts for intelligence analysts under data overload. In *Proceedings of the Human Factors and Ergonomics Society 45th Annual Meeting*, October 2001.

[10] P. Pirolli and S. Card. The sensemaking process and leverage points for analyst technology as identified through cognitive task analysis. In *Proc. of the International Conf. on Intelligence Analysis*, 2005.

[11] J. Scholtz, E. Morse, and T. Hewett. In depth observational studies of professional intelligence analysts. In *Presented at Human Performance, Situation Awareness, and Automation (HPSAA)*, 2004.