AFRL-RI-RS-TR-2009-193
Final Technical Report
August 2009

ENABLING A COLLABORATIVE PROBLEM-SOLVING FRAMEWORK THROUGH USER INTENT MODELING OF THE ANALYTIC PROCESS

Dartmouth College and University of Wisconsin-Whitewater (UWW)

***APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED*.**

STINFO COPY

AIR FORCE RESEARCH LABORATORY
INFORMATION DIRECTORATE
ROME RESEARCH SITE
ROME, NEW YORK

NOTICE AND SIGNATURE PAGE

AFRL-RI-RS-TR-2009-193 HAS BEEN REVIEWED AND IS APPROVED FOR PUBLICATION IN ACCORDANCE WITH ASSIGNED DISTRIBUTION STATEMENT.

FOR THE DIRECTOR:

/s/	/s/
JAMES M. NAGY	JOSEPH CAMERA, Chief
Work Unit Manager	Information & Intelligence Exploitation Division
	Information Directorate

REPORT DOCUMENTATION PAGE	*Form Approved* **OMB No. 0704-0188**

1. REPORT DATE *(DD-MM-YYYY)*	**2. REPORT TYPE**	**3. DATES COVERED** *(From - To)*
AUGUST 2009	Final	August 2007 – December 2008

4. TITLE AND SUBTITLE

ENABLING A COLLABORATIVE PROBLEM-SOLVING FRAMEWORK THROUGH USER INTENT MODELING OF THE ANALYTIC PROCESS

5a. CONTRACT NUMBER N/A

5b. GRANT NUMBER FA8750-07-2-0219

5c. PROGRAM ELEMENT NUMBER N/A

6. AUTHOR(S)

Eugene Santos and Hien Nguyen

5d. PROJECT NUMBER CASE

5e. TASK NUMBER BA

5f. WORK UNIT NUMBER 74

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

Dartmouth College
11 Rope Ferry Road, #6210
Hanover, NH 03755-1404

8. PERFORMING ORGANIZATION REPORT NUMBER

N/A

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

AFRL/RIED
525 Brooks Road
Rome NY 13441-4505

10. SPONSOR/MONITOR'S ACRONYM(S)
N/A

11. SPONSORING/MONITORING AGENCY REPORT NUMBER
AFRL-RI-RS-TR-2009-193

12. DISTRIBUTION AVAILABILITY STATEMENT

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED. PA# 88ABW-2009-3576

13. SUPPLEMENTARY NOTES

14. ABSTRACT
The research team during the CASE project was working on the problem of modeling an analyst's intent in order to improve tacit collaboration among analysts. The actions of an analyst are captured, and their goals and commitment to achieve them are inferred in order to improve effectiveness in collaborative tasks by building a user model to predict an analyst's intent. The targeted sharing of knowledge among analysts assists to reduce the time needed for junior analysts to solve problems by allowing the sharing and reuse of common/expert knowledge and actions that experts had in similar situations. This research helps ensure successful collaboration by development of analyst intent to provide deep understanding of what analysts are trying to achieve and how they are achieving their goals. Without this level of understanding of the analyst's intentions, collaboration may only be suggested based on similarity of topical interests among analysts which does not reflect the motivation behind these interests. This report is an overview of the state of the art research; the main idea of the researchers' approach; presents a scenario; and describes the valuable lessons learned from this project.

15. SUBJECT TERMS
Modeling Analytic Process; Tacit Collaboration; Prior knowledge; Hypothesis tracking; User modeling; Collaboration and Analyst/System Effectiveness; CASE

16. SECURITY CLASSIFICATION OF:			**17. LIMITATION OF ABSTRACT**	**18. NUMBER OF PAGES**	**19a. NAME OF RESPONSIBLE PERSON**
					James M. Nagy
a. REPORT	**b. ABSTRACT**	**c. THIS PAGE**			**19b. TELEPHONE NUMBER** (*Include area code)*
U	U	U	UU	25	N/A

TABLE OF CONTENTS

LIST OF FIGURES

1. INTRODUCTION

This report describes the technical and scientific results that our Dartmouth/ University of Wisconsin-Whitewater (UWW) team has accomplished during the scope of Collaboration and Analyst/System Effectiveness (CASE) project from August 2007 to August 2008. We are working on the problem of modeling an analyst's *intent* to improve the effectiveness of collaboration among intelligence analysts. Our approach offers a way to improve the *diversity* in a collaborative group by looking at the commonalities of the overarching goals that the analysts share instead of specific topics. Most of the existing approaches to modeling users for group collaboration explore the *similarity* of the users' topical interests (McDonald et al, 2000, Schmitt et al, 2003). There are two problems with this approach. First, people with similar interests may get stuck at the same peaks because they view and solve problems similarly (Page 2007). Real world collaboration will be beneficial from a group of people with different interests and perspectives who are working towards the same goal, as shown in many examples from various fields ranging from business to science in (Page, 2007). Secondly, topical interests only shows *what* the users have in common but do not show *how* the users are achieving or using these interests for their tasks.

We address these gaps by capturing the user's intent in which the intent is defined as his goals, commitment to achieve these goals, and actions leading toward his goals. We believe that with this level of understanding of the analyst's intentions, collaboration groups may be better formed with people who are working toward the same big goals and different courses of actions. Moreover, to improve the effectiveness of collaboration, it is crucial to find people with precise descriptions of their overarching goals and find them early enough to make the collaboration a success. Lastly, by capturing analyst intent, we will enable the targeted sharing of knowledge among analysts as well as potentially reduce the time needed for junior analysts to solve a problem by allowing the sharing and reuse of common/expert knowledge and possible goals and actions that experts had in similar situations.

This problem is particularly interesting and challenging in both the research and testing phases. It involves many open-ended research questions that are important for cross-disciplinary areas including information retrieval, goal recognition, and user modeling. This is particularly challenging from the testing point of view because it involves a number of issues in terms of assessment plans, scalability, and robustness across domains. During the 1 year period, our team has worked really hard to develop a framework for a new intent model, develop some main modules as web services as well as worked out a feasible testing plan to evaluate this model. In this document, we start first by describing our methods. This is followed by scientific results and lessons learned from this project. We also conclude by identifying our ongoing/future research direction.

2. METHODS AND PROCEDURES

Developing an intent-based user model is neither a new nor an easy problem. The difficulty of this problem is that it involves principles from several disciplines such as social science, computer science, and psychology. Additionally, the vagueness and dynamic nature of a user's behaviors poses a great challenge to designers of such a model. Our approach **differs** from existing approaches in two key aspects: (i) **dynamics** in building a goal hierarchy and action/goal relationships; and, (ii) providing information about **process** as opposed to a topical description of intent. By a goal hierarchy, we mean a hierarchical structure that describes the road leading up to achievement or accomplishment of a task for a user. During the one year period, we have successfully developed a formal framework for a new intent model, developed a scenario that shows our model will work using the APEX (National Institute of Standards and Technology (NIST) test bed of analyst logging experiment) dataset, and worked on an evaluation plan to assess this model.

2.1 Intent-based model framework

Our research on capturing a user's intent was spanned all the way back to 1999 (Santos et al, 1999, 2001, 2003a, 2003b, Nguyen et al., 2004a, 2004b, Nguyen, 2005; Santos and Nguyen 2008). Our definition of intent is also consistent with those found in the social sciences (Malle et al. 2003). We have developed a computational model to capture analyst intent by analyzing the actions taken by analysts as well as the contents of relevant documents/snippets/annotations arising from their actions. Our model uses dynamically constructed probabilistic, directed acyclic graphs, constructed in part by utilizing natural language processing techniques, to represent a user's intent. This model needs to provide the information on *What* the analyst's focus is (his goal); *How* committed he is to a particular goal; *How* the analyst is achieving this goal; and, *Why* the analyst is trying to achieve this goal. To answer the "Why" question, we use a *Rationale network* in which the relationships among the goals, information, and background knowledge triggered by a user's actions are stored. We also use an *Action network* which captures the relationships between goals and sequences of actions leading toward these goals which can thus be used to provide information for the "How" question. We use a *Foci network* to capture an analyst's current foci (goals) which can be used to provide information on the "What" question. Also, in this network, we capture the level of commitment of an analyst to a particular goal by measuring the frequency of how many times a focus is directly or indirectly pursued, and how recently a focus was pursued (e.g., abandoning information).

2.1.1 Definition

Intent is a tuple $I = \{G,A,C\}$ in which G is a goal or a set of goals, A is a set of actions to achieve this goal(s), and C is a real value(s) indicating how committed an analyst is to this goal(s). A goal is an end-state that an analyst is trying to achieve while solving a task at hand. For example: the analyst's goal could be to successfully justify in his report that Iranian leaders support a nuclear weapons program. We classify goals based on their persistence and content. Specifically, we classify goals into long-term and short-term goals. Long-term goals are abstract concepts, or general, high level hypotheses being explored while short-terms goals are specific concepts and triggered by low-level actions. In order to achieve a goal, analysts perform a set of actions, e.g. searching for relevant documents on a specific topic and annotating relevant documents to describe important points. For the CASE project, actions can be described by an Analysis Log Event (ALE). An action is pertinent to a goal if (i) it directly triggers the goal; or (ii) it triggers another goal that is a sub-goal of the goal. A goal is instantiated if it is created by goal-directed actions; or, if the goal-directed actions and pre-conditions of the goal are matched with past information and actions. A goal is satisfied if the analyst has explicitly specified it is, or if the analyst has collected information about all related concepts of that goal. A goal is abandoned if the user's actions have not triggered this goal after a fixed amount of time that can be specified by the analyst.

2.1.2 Components

2.1.2.1 Rationale network

A rationale network is a directed acyclic graph that consists of 4 types of *nodes*: (i) *Context*: is directed acyclic graph that includes concepts and their relations that are extracted from the content of documents, snippets, annotations that an analyst is using writing, copying, annotating and provided by other CASE teams; (ii) *Beliefs*: what the analyst believes about something or in something based only on collected/gathered info, e.g. sourcing evidence (bottom-up, collected information); (iii) *Goals*: what the analyst is aiming for or trying to reach/prove; and, (iv) *Axioms*: what the analyst believes in *not* based on collected info, e.g. intelligence doctrine/training of the analyst – bias can be one aspect of the axioms. There are four types of links

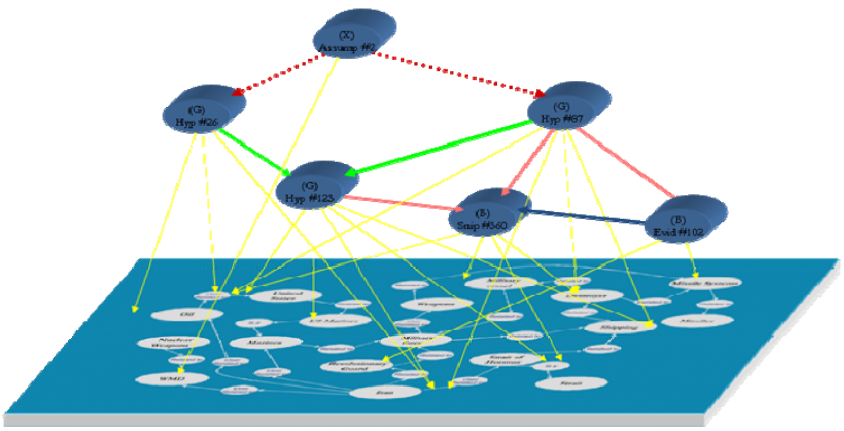

Figure 1: An example of a rationale network

in the rationale network: (i) *Context Connections:* the relationships between concepts that are the basis for goal, belief, and axiom nodes; (ii) *Support Links:* the relationships between goal and belief nodes as well as goal and sub-goal nodes; (iii) *Motivation Links:* how individual knowledge (not obtained through collection activities) represented in axiom nodes impact belief and goal nodes; and, (iv) *"Links-To" Links:* heuristic connection of nodes based on observations of the analyst, e.g., snippets occurring very close in time may be linked.

We connect all the concept nodes generated to the chosen goal. A belief node is added if a user makes explicit what he believes in (e.g., statements in an annotation). A new goal node should be created if the set of belief and context nodes of the existing goal is only covered by that goal at most *t*% of the time with *t* being the cutoff threshold. The name and description of the new goal node will be generated from the set of context, axiom and belief nodes.

2.1.2.2 Action network

An action network is a Bayesian knowledge-base (Santos & Santos 1999) that contains 2 types of nodes: goal and action nodes, and satisfies the following constraints: (i) goal nodes can be connected to action nodes or other goal nodes; and, (ii) action nodes can be connected to other action nodes. An action network is created from the user observed actions (ALE-based actions), goals from the rationale network and a Hidden Markov Model (HMM) initially with 3 states to determine the next possible user action. Once a user's behavior is reported and retrieved (e.g., using the ALE format), a context graph related to that behavior is generated. A possible action will be predicted based on HMM module. A goal node(s) with the highest value from the rationale network is obtained by doing reasoning over the Rationale network (G_r). The set of current actions, goals, and predicted next actions will be used to predict the next goal on the action network (G_a). We will take the intersection between G_r and G_a. If this intersection is empty then we will create a new goal node based on the information retrieved of these two sets. An edge is drawn between the goal node and the action node. The probability of the action node given this goal node is determined based on the weight of the goal node in the rationale network.

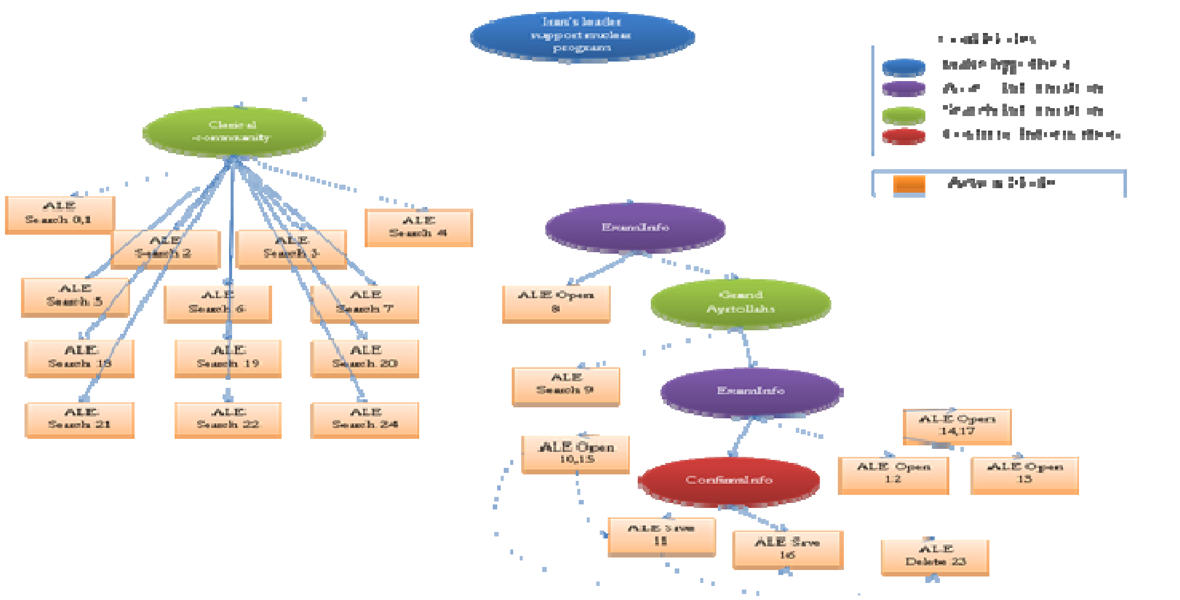

Figure 2: An example of an action network

2.1.2.3 Foci network

The Foci network is also a directed acyclic graph *(V, E)* in which *V* is a set of goal nodes that represents the current foci of an analyst and *E* is a set of edges that connect the goal nodes. Each node is described by a goal node, a set of weighted interests that are relevant to this goal, and a real number representing the commitment level for the goal, currently computed from the frequency and recency of the goal being pursued. Each edge is described by the source and destination goal nodes and the type of link it represents. There are two types of links: *regular* links and *leakage* links. A regular link represents the link between two goals as it is shown in the Rationale network while a leakage link represents a relationship in which two goals have been fired together frequently. In the foci network the goals are partitioned into a set of long-term goals and a set of short-term goals.

A new goal node is created when the current action is not consistent with the current top N goal nodes activated in the rationale network and it does not exist in the existing foci network. The set of affiliated interests/context is extracted from the document related to this action (e.g., from an annotation, a snippet, or from topics of a downloaded web page). The level of commitment is initialized to the goal value in the rationale network if this is a new goal. The edges coming into/out of this newly added goal node need to be consistent with the hierarchy in the rationale network. With the potential addition of new focus, the foci network is updated using spreading activation (Anderson 83) to propagate activity levels throughout the network with the modification that: (i) propagation is stronger along white arrows; (ii) propagation is weaker along leakages; and, (iii) propagation is universally stronger for arrows to/from a long term focus vs. to/from a short term focus. A goal node can become long term if the analyst indicates it is explicitly – e.g. through long-term task descriptions, etc., or if the depth of its descendants surpasses a given threshold (long term hierarchy threshold).

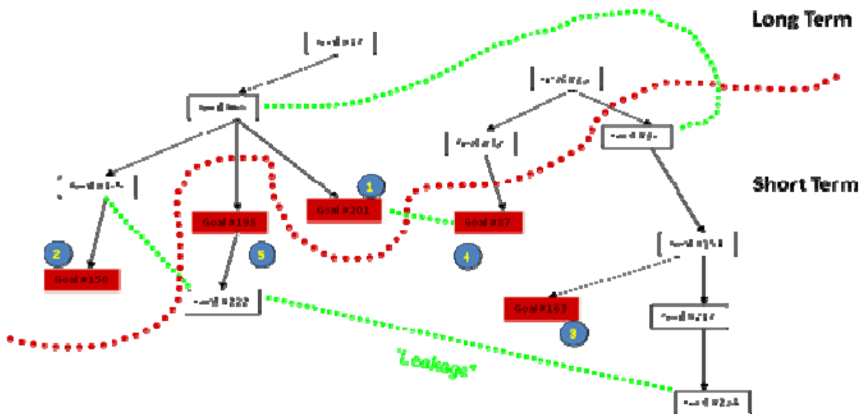

Figure 3: An example of a foci network

Intent is thus determined by first setting the currently active foci and those context nodes involved in the currently observed action in the Rationale network as evidence; then propagating interest through the Rationale network using spreading activation to infer the most active goals.

2.2 Scenario

In this scenario, we show how intent is captured and presented. We use the APEX collection in developing this scenario. The APEX collection, offered by NIST, consists of 8 analysts. This collection included 8 analysts, their recorded actions over time, and their research reports and assessment reports generated on their analysis. Each analyst was requested to assess the two hypotheses: “Where does the Iranian clerical community stand on Ayatollah Khamenei and President Ahmadinejad’s policies with regards to Iran’s civilian and military nuclear program?” Their actions are captured and stored in a common repository using Analysis Log Event (ALE) format (described in CASE Analysis Log Service Specification). Each ALE contains information about analyst's name, the time when the action took place, type of action, and content of an action for some selected actions.

We chose 2 users from APEX collection: APEXB and APEXF because they are very similar yet contain distinct differences. They are similar because they are always in the same cluster if we cluster these 8 analysts using K-means algorithm on research, assessment reports, the sets of queries, and visited documents. However, based on their assessment reports, their own final conclusions are still distinct (APEXB assessed that there are fissures in the clerical community but they do not present a deepening divide among the clerics loyal to the Iran revolution while APEXF assessed that the clerics actively oppose the regime). This makes them a good pair for collaboration because they have different perspectives.

In this scenario, we took a sequence of actions for each user, construct a user model for each user based on these actions and infer each user’s intent at a given time. Figures 4 and 5 show the sequences of actions of APEXB and APEXF, extracted from the set of actions that they performed and are recorded in the APEX collection.

ALE

1. Search: "who are Clerics who support Iran's conservative government"
2. Search: "where are debates between clerics"
3. Retain: 495cce1bX11635035eb9XY32a4
4. Search: How many debates do you find between Iranian clerical leaders
5. Search: who is Mahmoud Ahmadinejad
6. Search: who is Ayatollah Ali Khamemei
7. Retain: 495cce1bX11635035eb9XY85a
8. Search: what are Debates between Iranian Clerics
9. Retain:495cce1bX11635035eb9XY6730
10. Retain:495cce1bX11635035eb9XY4afc
11. Write to research report: "Supreme Leader Ali Khamenei gave his unequivocal backing to the country's nuclear program, saying it was the future and destiny of his regime to continue with the program ."

Figure 4: APEXB’s sequence of action

1. Search: Who are the Grand Ayatollahs in Iran?
2. Retain: 495cce1bX11635035eb9XY3b14
3. Search: how do clerics show their attitude toward the current conservative government?
4. Retain: 495cce1bX11635035eb9XY5042
5. Search: what are Iranian cleric alliances against hard-line government policies
6. Retain: 495cce1bX11635035eb9XY3b14

Figure 5: APEXF's sequence of actions

The numbers shown in the sequences of actions and figure labels denote the point in time that the actions took place. For example, at time 1, APEXB issues a query about the Iranian clerics who support Iran's conservative government. We construct a user's Foci, Rationale and Action networks from these actions. An example of the Foci network for APEXB at time 3 is shown in Figure 6. At time 3, the current goal of the user APEXB is to find out information about Iranian clerics and nuclear weapons. This goal is further strengthened at time 11 as shown in Figure 7.

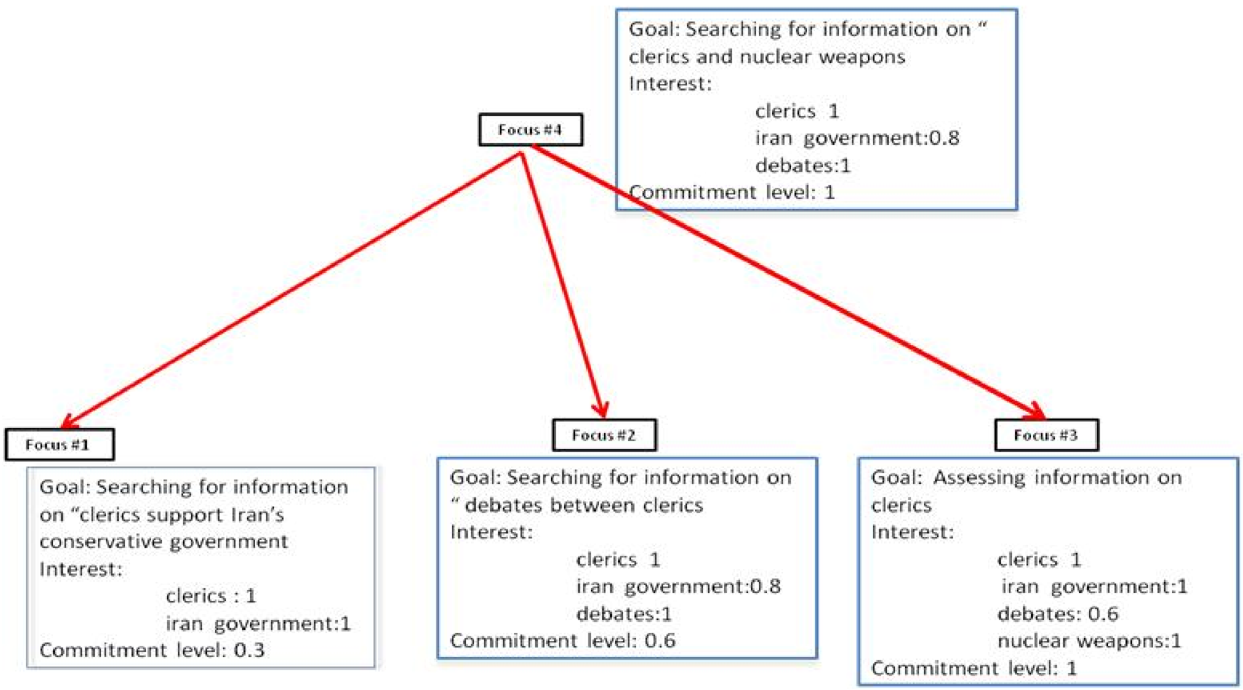

Figure 6: Foci network at time 3

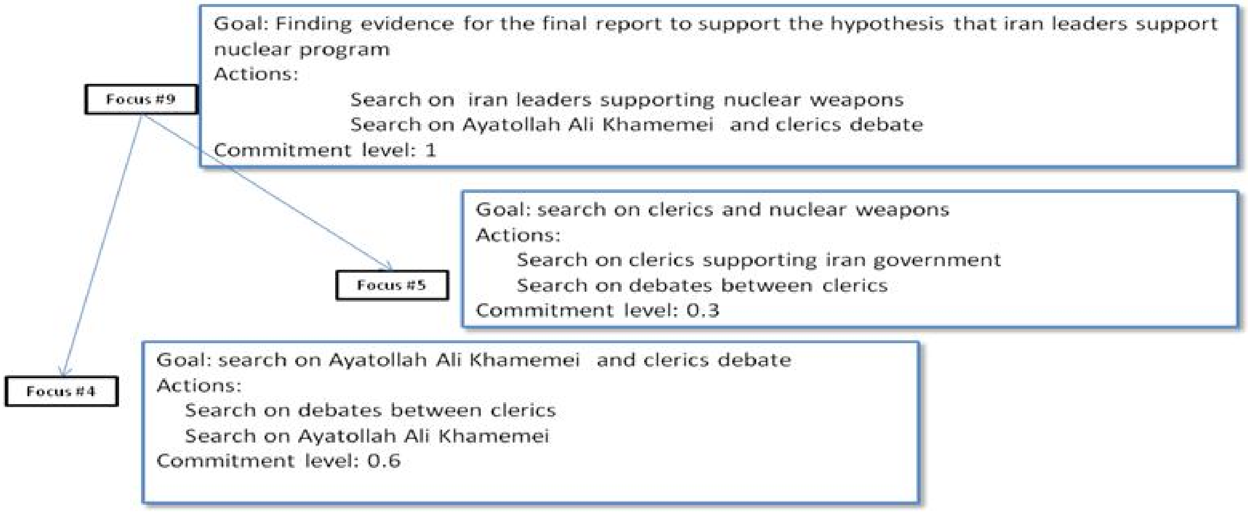

Figure 7: APEXB's intent at time 11

Figure 8 shows APEXF at time 6, our model suggests that the current intent of this user is to find information about Iranian leaders who support nuclear weapons. So the collaboration can be established very early in the analytic process despite the fact that the users are taking two different routes to achieve their goals, as well as reaching different conclusions at the end.

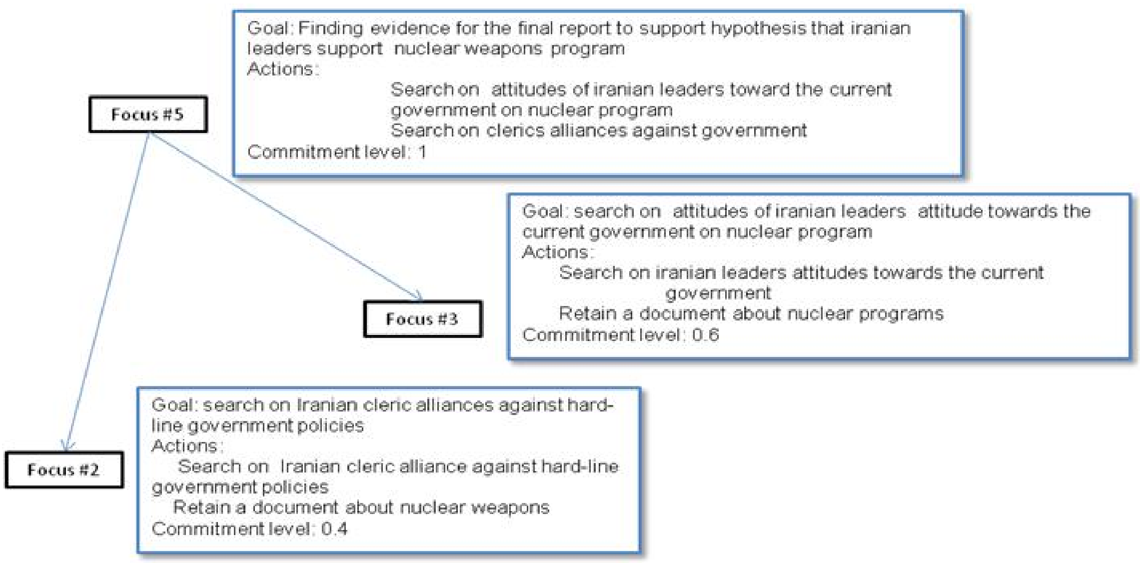

Figure 8: APEXF's intent at time 6

2.3 *Implementation*

2.3.1 Preliminary assessment of new intent model

Our objectives are to show that (i) we capture user intent more precisely in the analytical process compared to the simple interest lists; and (ii) we capture user intent *earlier* in the analytic process compared to the interest-based approach. These objectives help us to get closer to our ultimate goal which is to improve the *diversity* in a collaborative group by looking at the commonalities of the overarching goals shared by intelligence analysts.

For the first objective, we choose four pairs of analysts who have different actions (APEXL and APEXC, APEXE and APEXH, APEXL and APEXK, APEXF and APEXB). The intuition behind this selection is that it addresses the diversity issue by combining people with different actions because they offer different perspectives. We considered Retain and Search events in this experiment. These analysts have different actions because they always belong to different clusters when we use K-means clustering algorithm to cluster their set of queries. Additionally, even though they have the same overarching goals, their final reports have distinct conclusions.

In our first experiment, we ran our intent model 7 times. Each time, we used 25 consecutive events from each of the chosen analysts that represented the actions that the analyst has done on December 11, 2007. For each pair of analysts, we defined the precision of our intent model as the ratio between the number of relevant common goals of the two analysts in the pair over the number of common goals. A *common goal* is a goal node that is found in both intent models representing these corresponding analysts. For the interest model, we considered a set of common concepts found in both the interest lists as the set of common goals. We took the set of terms from the two working hypotheses as the ground truth of the analysts' goals. The average of precision for the interest model for these four pairs is 0.43 (sd=0.08), and for the intent model is 0.74 (sd=0.15). The paired t test results reveals that the results are statistically significant (n=4, p-value= 0.0396).

In the second experiment, we measured the time at which the common goals of these two analysts were found for our intent model and the model containing only interests. We chose APEXF and APEXB for this experiment. For each analyst, we created our intent model on the fly with the inputs from the set of 40 events and output three components of our intent model for each time slice. We chose 40 events for each analyst (APEXB and APEXF) on December 11, 2007 such that they did not start with the same focus. APEXB started with the question on ``nuclear weapon program and Iran" while APEXF asked about “grand Ayatollah". We found out that at time *t=5*, our intent model has precisely picked up the common goals of Iran nuclear program and cleric leaders while at time *t=8*, the interest model has picked up ``cleric", ``Iran", ``nuclear" as interests.

This experiment gives us some insights to develop a more comprehensive evaluation plan in which we divide the set of events for each analyst into a set of sessions and perform similar assessments over the numerations of the set of sessions of all analysts.

2.3.2 Existing web service user model

From August 2007 to August 2008, we supplied a service called *SupplyUMInfo* that accepts an analyst's name as input and returns the concepts of interest to the analyst along with the level of interest in each concept. Specifically, *SupperUMInfo* is deployed as web service using Tomcat Apache Web server, AXIS 1.4 at the University of Wisconsin-Whitewater. (http://140.146.84.182/axis/services/SupplyUMInfo?wsdl). BAE team has used this web service in their Expert profile page (shown in Figure 9). All data shown is fictitious and the red box would illustrate picture of the expert.

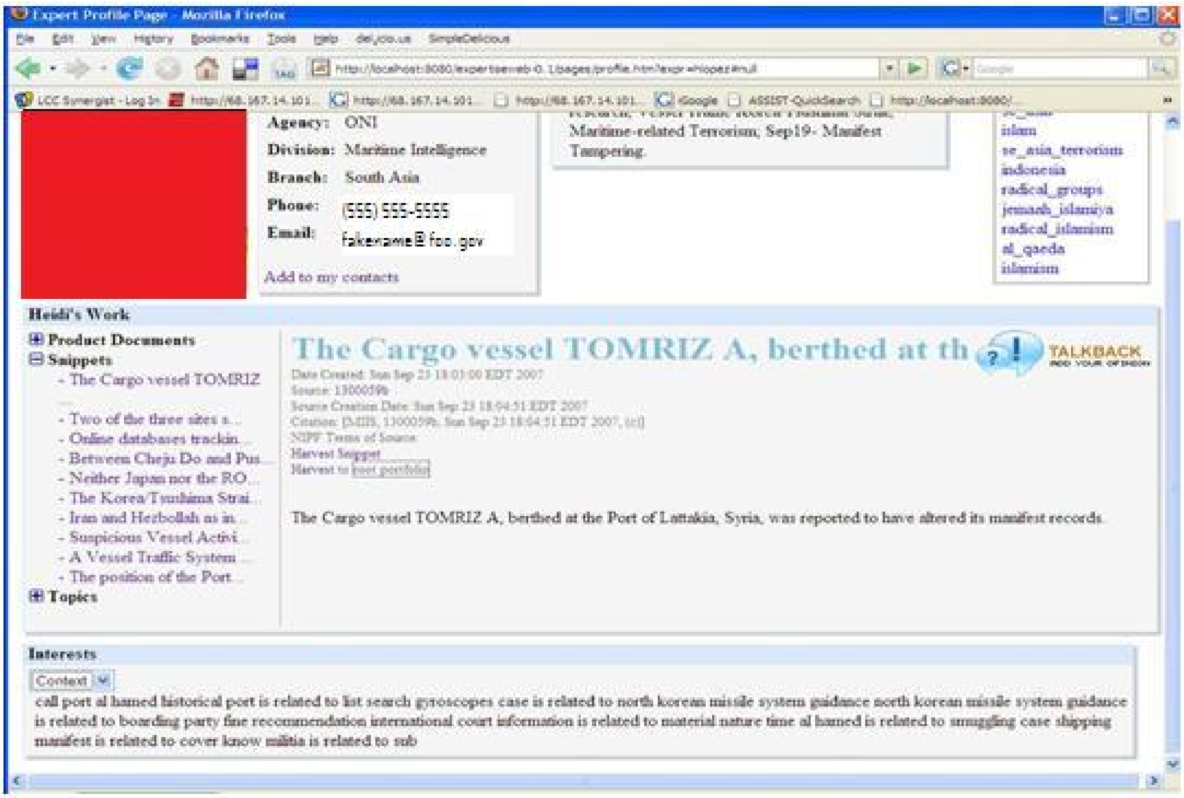

Figure 9: SuppyUMInfo service used in BAE Expert profile page

2.4 Vision of our project

Analysts face tremendous challenges on a daily basis that include gathering relevant information quickly and accurately from massive dynamic information spaces; performing precise analyses to identify and connect evidence to hypotheses; and, providing immediate interpretation on certain events based on said hypotheses, all to solve just one of their many tasks at hand. Many techniques have been developed to help analysts such as information filtering agents (Billsus & Pazzani, 2000), (topically) clustering the search space (Melvin, 2004; Wang & McCallum, 2006), using predefined taxonomies (Tanaka et al, 1999), or determining prior and tacit knowledge (Cheng et al., 2005). While these techniques have demonstrated much potential and value in assisting the analyst, they only scratch the surface of a deeper fundamental requirement, that of accounting for and understanding the needs of the analyst especially as unique individuals with their myriad of differences in style, ability, and operation. Hence, these tools run into the danger of tremendous information "push" without the critical attendant "pull" by the analyst since the tools have an incomplete account of the analyst.

With our intent modeling technique, we can assist the analyst in a number of ways ranging from helping to better identify relevant snippets of information and enhancing awareness of unknowns/biases, as well as pointing out potential collaborative analyst activities by analyzing and reasoning with in-depth the various user models; to employing user models of senior analysts to assist in the training and activities of junior analysts through suggestions and potential critiques. Furthermore, such an intent modeling approach will allow collaborating groups of analysts to achieve a specific overarching goal while improving communications between analysts via a common representation (inclusive of analyst contextual knowledge) and while actively providing assistance, as needed. It helps improve "diversity" in group collaboration which, unfortunately, has been overlooked in the existing approaches to create a collaborative group.

3. SCIENTIFIC RESULTS

3.1 Published Papers

There have been published one book chapter and two conference papers from the work related to this project.

3.1.1 Book chapter(s)

- Modeling Users for Adaptive Information Retrieval by Capturing User Intent. Eugene Santos Jr. Hien Nguyen. 2009. *Collaborative and Social Information Retrieval and Access: Techniques for Improved User Modeling*. IGI Global. (Submitted March, 2008, accepted July 2008. (Tentative schedule) Available for public in February 2009).

This book chapter described in detail our existing user model IPC which we used to implement the service *SupplyUMInfo* that accepts analyst's name as input and return the concepts of interests to the analysts along with level of interest for each concepts as output. This is a peer-reviewed book in which each chapter has gone through 2 blind review iterations in which a chapter is reviewed by 3 reviewers.

3.1.2 Conference paper(s)

- Eugene Santos Jr. Hien Nguyen, Fei Yu, Keumjoo Kim, Deqing Li, John Wilkinson, Adam Olson, Jacob Russell. Capturing User Intent for Analytic Process. 2009. In *Proceedings of 2009 User Modeling Adaptation and Personalization.* Trento, Italy.

The above paper presents our new intent model as described in this report as well as the results of our preliminary experiments.

- Intent-driven insider threat detection in intelligence analyses. 2008. Eugene Santos Jr. Hien Nguyen, Fei Yu, Keumjoo Kim, Deqing Li, John Wilkinson, Adam Olson, Jacob Russell. In *Proceedings of 2008 IEEE/WIC/ACM International Conference on Intelligent Agent Technology*, Sydney, Australia.

In the above paper, we develop a unified framework for intent-driven insider threat detection. The heart of the framework is the IPC user modeling technique which captures the analyst's interests, knowledge context, and preferences over time. We conducted an empirical evaluation using the APEX '07 collection. The APEX dataset was created by the National Institute of Standards and Technology (NIST) to simulate an analysis task in the intelligence community. The APEX '07 collection included 8 analysts, their recorded actions over time, and their research reports as well as assessment reports generated on their analysis. Five malicious insiders were simulated each based off of one of the original 8 analysts. We measured the similarities between the final user model and different hypotheses in the assessment reports for all analysts. In order to analyze these similarity values, three different metrics were proposed to compare the deviation values between multiple hypotheses either in each section of the assessment report or in the entire report for identifying suspicious insiders. The experimental results showed that the framework was effective in identifying insider threats. The first and third metrics detected four insiders with malicious intent. The third metric did not raise any false positives while the first and second metrics had false positives on two benign analysts.

3.1.3 Related papers

- Evaluation of the Effects of User-Sensitivity on Text Summarization.2008. Hien Nguyen, Eugene Santos Jr., Russell Jacob, and Nathan Smith. In *Proceedings of 2008 IEEE/WIC/ACM International Conference on Web Intelligence.* Sydney, Australia.

- Impacts of User Modeling on Personalization of Information Retrieval: An evaluation with hyman intelligence analysts. Eugene Santos Jr., Qunhua Zhao, Hien Nguyen, Hua Wang. 2005. In *Technical report of Workshop on Evaluation of Adaptive Systems at UM 2005.*

3.2 Technology transition status

Technology delivered to intelligence community organization/program: our SupplyUMInfo service that uses ALS/ALE has been deployed and used by BAE system in their product. We plan to deploy our new intent model in the same way. Also we plan to apply this technique in our framework for intelligent foraging, gathering and matching (IFGM) (Santos et al., 2008).

4. LESSONS LEARNED

There are several valuable lessons that we have learned from participating in this project regarding basic research, integration, and evaluation.

4.1 Basic research

The most challenging questions are to quantify an analyst's intent, and capture relevant information to build a computational model of the analyst's intent. Capturing user intent is known to be a challenging problem spanning across multiple domains because intent modeling must deal with vagueness, dynamics, and partially-observable information. There is a big gap between theories about intent in sociology and computational models and representation techniques in artificial intelligence. Additionally, we need to define analyst's intent during the analytic process in the intelligence community. This application domain adds another level of difficulty to the problem because there is limited literature on the analytic process and restricted access to information in the intelligence community. We have addressed this problem by using definitions from sociology about a user's intent to determine what is relevant and measurable in this domain, our evaluations with human intelligence analysts in the past to determine the most important components, and the data collections available in CASE program to develop our own model. The valuable lesson that we learn from this process is that the data sets available in such a restricted domain can help us considerably in revising our initial research questions so that we can define the intent model for this domain more precisely.

4.2 Integration

The most useful lesson that we learn from doing integration with other teams is the usefulness of Analytic Log Services (ALS) and Analytic Log Event (ALE) and the convenience of web service paradigm. ALS/ALE and the web service paradigm give us the data and mechanism to work independently and simultaneously with other teams so that we can start our project and integrate with other teams very quickly. Two weeks after the project officially started, our team was able to deploy *SupplyUMInfo* services which used ALEs generated from other teams, especially BAE systems. Without the use of ALS/ALE and the web service technology, it would definitely take longer time.

Additionally, the challenging problem for integration that we learned is to define protocols to communicate with other teams within the CASE program in order to best get the information for our model and to best serve other team's interests. We learned that by keeping other teams informed frequently on our goals as well as paying attention to detailed technical documents, we can serve/communicate with other teams better. We learn that in order to use and to provide web services, we need to check what come from other sites as well as to ensure that the servers and related components are on and running all the time at our site.

4.3 Evaluation

We learned that having data collections to work with is an invaluable tool for our research brainstorming as well as implementation. The set of ALEs from other teams, especially from BAE systems, helps us evaluate each component in our user model service and the integration of those components very quickly. Thanks to the CASE program, we have access to the Monterrey and the APEX collection which have helped us tremendously in both the design of our new model as well as in planning for the evaluation of the model.

We also learned that there are many other research questions that can be addressed using our user model on the data sets such as APEX and Monterrey collections. The work on detecting malicious insider by modifying the analysts from APEX collection (published in Intelligent Agent Technology conference 2008) is just one example.

5. ONGOING AND FUTURE WORK

In this section, we describe the future directions that our team will take regarding the implementation as well as research questions. First, we have now finalized our development of an intent-based user model that is based on what we have described above. We are done with the design for each unit and completed the implementation of each component of this model by the end of October 2008. Secondly, regarding the basic research questions, we would like to extend the current model and focus on developing an evaluation instrument and more test beds and scenarios to evaluate this model. The current model infers a user's intent primarily based on the documents that users read, annotate, print, and bookmark. We would like to extend to accommodate the other types of inputs that aren't documents, such as the situations when the tasks are given, the prior knowledge of the users and so forth. Also the evaluation plan right now is made under the assumption that we have somehow access to the pool of human analysis. This plan may not be feasible without the help of NIST. Therefore, we need to look at other similar domains such as business or finance analysts and customize our tasks and evaluation accordingly.

6. REFERENCES

- Anderson, J. R. (1983). A spreading activation theory of memory. *Journal of Verbal Learning and Verbal Behavior,* 22, 261-295.
- Billsus, D. and Pazzani, M. (2000). User Modeling for Adaptive News Access. *UMUAI* 10(2/3). 147-180.
- Cheng, J., Emami, R., Kerschberg, L., Santos, Jr., E., Zhao, Q., Nguyen, H., Wang, H., Huhns, M., Valtorta, M., Dang, J., Goradia, H., Huang, J., Xi., S. (2005). OmniSeer: A Cognitive Framework for User Modeling, Reuse of Prior and Tacit Knowledge, and Collaborative Knowledge Services. *Proceedings of HICSS-38*.
- Malle B, Moses, L., Baldwin D. 2003. *Intentions and intentionality: Foundations of Social Cognition*.
- McDoland, D. W., and Ackerman, M. S. 2000. Expertise Recommender: A plexible recommendation system and architecture. In *Proceedings of the 2000 Conference on Computer-Supported Cooperative Work.* pp 231--240. Philadelphia, PA.
- Melvyn, S. (2004). Semi-Automatic Taxonomy for Efficient Information Searching. *Proceedings of the 2nd International Conference on Information Technology for Application*.
- Nguyen, H. 2005. Capturing User Intent for Information Retrieval. Ph.D dissertation. University of Connecticut.
- Nguyen, H., Santos E. Jr., Zhao, Q., Lee, C. 2004a. Evaluation of Effects on Retrieval Performance for an Adaptive User Model. In: *Adaptive Hypermedia 2004: Workshop Proceedings -Part I,* (pp. 193–202), Eindhoven, the Netherlands
- Nguyen, H., Santos, E.J., Zhao, Q. & Wang, H. (2004b). Capturing User Intent for Information Retrieval. In: *Proceedings of the Human Factors and Ergonomics society 48th annual meeting*. (pp. 371–375), New Orleans, LA.
- Page, S., E. 2007. The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies. *Princeton University Press*.
- Santos, E. Jr, Nguyen, H., Zhao, Q. & Pukinskis, E. (2003b). Empirical Evaluation of Adaptive User Modeling in a Medical Information Retrieval Application. In: *Proceedings of the ninth User Modeling Conference UM 2003*, (pp. 292–296). Johnstown. Pennsylvania.
- Santos, E. Jr, Nguyen, H., Zhao, Q., & Wang, H. (2003a). User Modeling for Intent Prediction in Information Analysis. In: *Proceedings of the 47th Annual Meeting for the Human Factors and Ergonomics Society (HFES-03)*, (pp. 1034–1038).
- Santos, E. Jr. and Santos, E. Sr. 1999. A Framework for Building Knowledge-Bases under Uncertainty. *Journal of Experimental and Theoretical Artificial Intelligence*, vol. 11, pp. 265-286.
- Santos, E. Jr., Brown, S. M., Lejter, M., Ngai, G., Bank, S., & Stytz, M. R. 1999. Dynamic User Model Construction with Bayesian Networks for Intelligent Information Queries. In: *Proceedings of the 12th International FLAIRS Conference*. pp. 3–7. Orlando. FL.
- Santos, E. Jr., Nguyen, H., & Brown, M., S. 2001. Kavanah: An active user interface Information Retrieval Application. In: *Proceedings of 2nd Asia-Pacific Conference on Intelligent Agent Technology*, (pp. 412–423).

- Santos, E. Jr., Santos, E., Nguyen, H., Pan, L., Korah, J., Xia, H., Yu, F., Li. D. 2008. Analyst-Ready Large Scale Real Time Information Retrieval Tool for E-Governance.(Chapter XVI). *E-Government Diffusion, Policy, and Impact: Advanced Issues and Practices*. Edited by Mehdi Khosrow-Pour. IGI Global. ISBN: 978-1-60566-130-8.
- Santos, E. Jr., Zhao, Q., Nguyen, H., Wang, H. 2005. Impacts of User Modeling on Personalization of Information Retrieval: An evaluation with human intelligence analysts. In S.Weibelzahl, A.Paramythis, and J. Masthoff (Eds.). *Proceedings of the Fourth Workshop on the Evaluation of Adaptive Systems*, held in conjunction with the 10th International Conference on User Modeling (UM'05), (pp 19-26).
- Schmitt, C., Dengler, D., and Bauer, M. 2003. Multivariate preference models and decision making with the MAUT Machine. In *P. Brusilovsky, A. Corbett, and F. de Rosis (Eds.), UM 2003, User Modeling: Proceedings of the Ninth International Conference. Berlin:Germany*, pp 148--159.
- Tanaka, H., Kumano, T., Uratani, N., and Ehara, T. (1999). An efficient document clustering algorithm and its application to a document browser. *Information Processing and Management* 35. 541-557.
- Wang, X. and McCallum, A. (2006). Topics over time: a non-Markov continuous-time model of topical trends. *Proc of the 12th ACM SIGKDD*. 424-433.

7. LIST OF SYMBOLS, ABBREVIATIONS AND ACRONYMS

- ALE – Analysis Log Event
- ALS – Analysis Logging Service
- APEX – NIST testbed of analyst logging experiment
- APEXx – Denotes the analysis x in APEX testbed
- BKB – Bayesian Knowledge Base
- CASE – Collaboration and Analyst/System Effectiveness
- HMM – Hidden Markov Model
- IFGM – Intelligent Foraging, Gathering, and Matching Information Retrieval System
- NIST – National Institute of Standards
- UM – User modeling
- UWW – University of Wisconsin at Whitewater