Capturing User Intent for Analytic Process

Eugene Santos Jr.[1], Hien Nguyen[2], John Wilkinson[1], Fei Yu[1], Deqing Li[1], Keum Kim[1], Jacob Russell[2], and Adam Olson[2]

[1] Thayer School of Engineering
Dartmouth College
8000 Cummings Hall Hanover, NH 03755
[2] Mathematical and Computer Sciences Dept.
University of Wisconsin Whitewater
800 W. Main Street
Whitewater, WI 53190

Abstract. We are working on the problem of modeling an analyst's intent in order to improve collaboration among intelligence analysts. Our approach is to infer the analyst's goals, commitment, and actions to improve the effectiveness of collaboration. This is a crucial problem to ensure successful collaboration because analyst intent provides a deeper understanding of what analysts are trying to achieve and how they are achieving their goals than simply modeling their interests. The novelty of our approach relies on modeling the process of committing to a goal as opposed to simply modeling topical interests. Additionally, we dynamically generate a goal hierarchy by exploring the relationships between concepts related to a goal. In this short paper, we present the formal framework of our intent model, and demonstrate how it is used to detect the common goals between analysts using the APEX dataset.

1 Introduction

We study 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 [6, 12]. 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 [7]. Secondly, topical interests only show *what* the users have in common but do not show *how* the users achieve or use these interests for their tasks. We address these gaps by taking the **first** step to capture the user's intent where the intent is defined as an analyst's goals, commitment to achieve these goals, and actions leading toward these 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 action. 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.

This problem is challenging because it involves several fields in the design and evaluation of an intent model, including sociology, computer science, and psychology. Two important research questions need to be addressed: (i) What is user intent and how do we capture it?; and (ii) How do we evaluate the effectiveness of the intent model? Our approach differs from existing approaches that capture a user's intent in an information seeking task (such as [1], [2], [4],[3], [9]) in that our model provides information about the *process* of a user's intent as opposed to a simple categorization of intent. This model is different from our previous user modeling approach ([10], [11]) in that the previous model focuses on capturing a user's topical interests as opposed to the process of achieving an overarching goal.

We developed a computational model to capture user intent by analyzing the actions taken by the user as well as the contents of relevant snippets and documents arising from his actions. Our model dynamically creates a goal hierarchy by finding the common concepts shared by directed acyclic graphs representing the relevant information. We capture the information on *What* the user's focus is (his goal), *How committed* he is to a particular goal, and *Which* actions he has taken to achieve this goal.

We demonstrate how our intent model is used to capture an analyst's intent by two simple experiments using the APEX dataset, which was created by the National Institute of Standards and Technology (NIST) to simulate an analytical task in the intelligence community. This collection included 8 analysts, their recorded actions over time, and their final reports. The preliminary assessment shows that our intent model captures the overarching goals more precisely and earlier in the analytical process than the model capturing only a user's interests. This paper is organized as follows: we describe our framework in detail. Next, we present two experiments with four pairs of analysts in the APEX collection. Finally, we present our future work.

2 Our Intent Model

Definition: We define a user's intent (I) as a tuple $I = \{G,\ A,\ C\}$ in which G is a set of goals, A is a set of actions to achieve these goals, and C is a set of real value(s) indicating how committed an analyst is to each goal in G. Our definition of intent is consistent with those found in the social sciences [5]. Our goals are characterized by their category and content. The category represents the user's intent generally, such as "Searching for evidence", "Going through a set of documents", while the content represents the detail information, such as "Imar's leaders support nuclear programs." Note that the names in this paper are changed. Based on our definition of intent above, the 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 finally, *Why* the analyst is trying to achieve this goal. Our aim is to explore the relationships among the components related to goal, actions and commitment to tie them together in the intent framework. Therefore, this model has three inter-related components: *Rationale*, *Foci* and *Action* networks.

Rationale network: A Rationale network is a directed acyclic graph (DAG) that consists of 2 types of nodes: (i) **Context**: includes concept and relation nodes that are extracted from the content of documents, snippets, annotations generated by an analyst; and (ii) **Goals**: represent what the analyst is aiming for. These goal nodes represent the detail information and are called content-based goal nodes. There are "context" links between context nodes, "support" links between context nodes and goal nodes, and "link-to" links between goal nodes. We construct the Rationale network from a user's query, and relevant snippets and documents as follows: (i) Convert a user's query, snippets or relevant documents into a document graph (DG) representation. The *DG* representation has been used in our prior work for building user models for information retrieval [10], [11]. "Context" links are created between these context nodes. (ii) Insert a content-based goal node into the Rationale network and add the "support" links from this goal node to all the concept nodes generated in Step (i). (iii) Update the Rationale network by finding the common ancestors of the concept nodes that are the children of the newly added goal node with the sets of concept nodes associated with the existing goals. If such an ancestor is found, a goal node is created and the link-to connections are created between the common ancestors and the existing goals. An example extracted from a Rationale network built for APEXF analyst in our experiment shows that the analyst focuses on a common goal of "nuclear program Imar", which are supported by two sub goals "Retain a snippet representing the Grand Aya Ali al-Sistani", and "Searching information on which Imarian clerical leaders debate". These two subgoals, in turns, are supported by context nodes such as "decision maker", "nuclear program", and "grand Aya".

Foci network: A Foci network is a snapshot of the Rationale network with additional information on commitment level and interest list. Each node has a name, a set of weighted interests, and a real number representing the commitment level for the focus. The name of a node in this network is the same as a name of a content-based goal in the rationale network. The set of interests consists of the context nodes which are the children of the corresponding content-based goal in the rationale network. The weight for each interest is the ratio of the frequency of the given interest concept over the total concepts related to the given goal. The commitment is currently computed by a linear function over the frequency and recency of the focus being pursued. The frequency is the ratio of the number of times this goal occurs in the rationale network over the total time slices. The recency is computed as follows: $(1 - (t - t_i))/(t + 1)$ in which t represents the current time slice and t_i represents the latest time slice this goal is active.

Action network: An action network has two components: a long-term component represented in a Hidden Markov Model (HMM) ([8]) and a short-term component represented in a Bayesian network. The HMM contains 3 states and

8 observations representing possible states and actions in an analytical process. The 3 states are "Searching for Evidence", "Going through documents", and "Examining evidence". The 8 observations are "Start application", "Search", "Retain" (triggered when an analyst bookmarks, prints, saves a document, or cuts and pastes information from a document to his/her report), "Access" (triggered when an analyst opens a document to view), "Make Hypothesis", "Associate Evidence" (triggered when analyst links a document or a snippet to a hypothesis), "Assess" (triggered when analyst assesses how relevant a document or snippet to a hypothesis), and "Discard" (triggered when a user discards evidence). The Bayesian network contains category-based goal and action nodes, and the links from category-based goals to actions. A category-based goal node is inferred from the HMM. We use a frequency table to update the conditional probability table for each node in the action network. An example extracted from an action network in one of our experiments shows that the analyst is searching for evidence and has taken several searches on "Imarian clerical community stand on Aya and president Amar's policies with regards to Imarian's civilian and military nuclear program", and "clerics who support Imarian nuclear program".

Intent inference: we determined the intent information as follows: (i) G is determined by finding the nodes in the Foci network with the highest commitment. Set them and their related context nodes in the Rationale network as evidence. (ii) A spreading activation process is performed on Rationale network to find the set of the most active goals. We added those goals to G. (iii) The action nodes that relate to these content-based goals with the corresponding time in the action network, are set as evidence. We perform a belief update and find the category-based goals in the action network with the highest marginal probability.

3 Preliminary assessment

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. We use the APEX collection (offered by NIST), which has 8 analysts. Each analyst was requested to assess the two hypotheses: "Where does the Imar clerical community stand on Aya?" and "President Amar's policies with regards to Imar's civilian and military nuclear program?". Their actions are captured and stored in a common repository. There are 5613 events in total.

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 Imar" while APEXF asked about "grand Aya". We found out that at time *t=5*, our intent model has precisely picked up the common goals of Imar nuclear program and cleric leaders while at time *t=8*, the interest model has picked up "cleric", "Imar", "nuclear" as interests.

This scenario 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.

4 Conclusion and Future work

In this paper, we have described the intent model that is used to capture a user's intent in an analytical process. The intent is defined as a set of goals that a user is trying to achieve, a set of actions leading toward the goals, and commitment level that represents how committed the user is to those goals. Our formal framework contains three inter-related components: Rationale, Foci, and Action networks. We develop two simple experiments in which we show that, by capturing the overarching goal of an analyst, it may help precisely describe what he is actually trying to achieve, comparing to listing a set of topics that he currently is focusing on.

There are many interesting and potential directions that we continue to address. In terms of implementation, the generation of a goal description from the set of information including content of relevant documents and query, descriptions of actions and description of the general goal of the analyst is needed to be coherent, logical and informative. We consider some heuristics to fuse several sources of information. In terms of evaluation, we look forward to extending beyond the development of the proof-of-concept scenarios to confirm if the results in our preliminary assessment hold for all analysts on a much more comprehensive evaluation. In addition, we continue to use the APEX dataset and measure how accurate the actions (or a sequence of actions) are predicted. In terms of effectiveness to forming collaboration, we need to define a measure to assess the diversity of a collaborative group and how diversity can improve the effectiveness of collaboration. We plan to find out whether the group consisting of analysts recommended by finding the common intent is more diverse than the group with analysts recommended by the existing approaches such as collaborative filtering, and content-based filtering.

References

1. Baeza-Yates, R., Calderon-Benavides, L. and Gonzalez-Caro, C.: 'The Intention Behind Web Queries'. In *Proceedings of Spire 2006*. Scotland. pp 98-109. (2006).
2. Broder, A.: 'A Taxonomy of Web Search'. *SIGIR Forum*. Vol 36(2). 3-10. (2002)
3. Jansen, B., Booth, D., and Spink, A.: 'Determining the User Intent of Web Search Engine Queries'. In *Proceedings of WWW' 07*. Canada. pp 1149-1150. (2007).
4. Lau, T. and Horvitz, E.: 'Patterns of search: analyzing and modeling Web query refinement'. In *Proceedings of UM' 99*. J. Kay (Eds.) Springer-Verlag New York. pp 119–128. (1999)
5. Malle, B., Moses, L., and Baldwin D. :'Intentions and intentionality: Foundations of Social Cognition'. *The MIT press*. (2003).
6. McDoland, D. W., and Ackerman, M. S.: 'Expertise Recommender: A plexible recommendation system and architecture'. In *Proceedings of the CSCW' 00*. pp 231–240. Philadelphia, PA. (2000).
7. Page, S., E. : 'The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies'. *Princeton University Press*. (2007)
8. Rabiner R. L. : 'A Tutorial on Hidden Markov Models and Selected Applications in Speech Recognition ' In *Proceedings of the IEEE*. Vol 77 (2). 257-286. (1989).
9. Rose, D. E. and Levinson, D.: 'Understanding User Goals in Web Search'. In *Proceedings of WWW' 04*. New York, NY, pp 13-19. (2004).
10. Santos, E. Jr., Nguyen, H., Zhao, Q., and Pukinskis, E.: 'Empirical Evaluation of Adaptive User Modeling in a Medical Information Retrieval Application'. In *Proceedings of UM' 03*. P. Brusilovsky, A. Corbett, and F. de Rosis (Eds.) Springer. pp 292-296. (2003).
11. Santos, E. Jr., and Nguyen, H.: 'Modeling Users for Adaptive Information Retrieval by Capturing User Intent'. *Collaborative and Social Information Retrieval and Access: Techniques for Improved User Modeling*. IGI Global. (2009).
12. Schmitt, C., Dengler, D., and Bauer, M.: Multivariate preference models and decision making with the MAUT Machine. In *Proceedings of UM' 03*. P. Brusilovsky, A. Corbett, and F. de Rosis (Eds.). Springer. pp 297–302. (2003)