RESIN: A

RESOURCE BOUNDED INFORMATION GATHERING SYSTEM FOR VISUAL ANALYTICS

PAST PROJECTS

Meta-Cognition in Coordinators

RESIN: A RESource bounded INformation Gathering System for Visual Analytics

WLAN Resource

Mathematical Analysis of Uncertainty Propagation in Agent Control

Safety in Multi-Agent Systems

PI: Anita Raja

RA: Jia Yue, Zhiyong Guo, Evan Liu

Collaborators: Professor Bill Ribarsky, UNCC, Professor Ashok Goel, Georgia Tech, Professor Jing Yang, UNCC

Sponsor: DHS

Knowledge gathering and investigative tasks in open environments are very complex because the problem-solving context is constantly evolving, and the data may be incomplete, unreliable and/or conflicting. These tasks are time critical and typically involve identifying and tracking multiple hypotheses; gathering evidence to validate the correct hypotheses and eliminating the incorrect ones. Visual analytics is the science of applying reasoning and analysis techniques to large, complex real-world data for problem solving using visualizations.

In RESIN, we designed and developed a mixed-initiative reasoning agent that will assist investigative analysts in foraging tasks and performing predictive analysis using blackboard-based reasoning, visualization and an intelligent user interface. The agent leverages sequential decision making and an AI blackboard system to support hypothesis tracking and validation in a highly uncertain environment. The reasoning process involves the sequential execution of multiple knowledge sources, gathering large amounts of evidence, reasoning about incomplete and contradictory information, and supporting hypothesis tracking. The agent is also equipped with the ability to adapt its processing to available resources, deadlines and their current problem-solving contexts. Moreover, one notable characteristic of this agent is that it is equipped with technologies to predict the impact of future events by first

determining the missing information in current events and then identifying the event occurrence trends based on the historical events captured in the global terrorism database.

Predictive Analytics in RESIN

The hypothesis we plan to validate is: Given a current event, it is possible to predict future events by first determining the missing information about the current event and then determining the event trends based on the historical events captured in the global terrorism database. We also want to identify the impact level of the future events. The specific problem solving process can be summarized as the following three phases:

Fig1: Predictive Problem

Phase1- Category Prediction (CP) : To predict the missing or unknown category information (group name) based on the current event.

Phase2- Event Occurrence Prediction (EOP) : To predict the event occurrence of a certain group as identified from phase1 using the time series analysis on the historical data.

Phase3- Event Impact Level prediction (EILP) : To predict the impact level of future events by the pre-defined rules and

exponential smoothing method based on the occurrence trends of that group from phase 2

RESIN's Reasoning Process

The RESIN agent consists of an AI Blackboard, a TÆM.S. task structure library, a Markov Decision Process (MDP)solver and heterogeneous knowledge sources (KSs). The AI Blackboard contains current goals and reasoning results from processing existing information, which includes raw data, various problem-solving states, and partial solutions. TÆM.S. is an abstraction of the low-level execution model and captures uncertainty in outcome distributions. The MDP is a probabilistic model, which captures the essence of sequential decision processes, and is used to compute policies that identify, track, and plan to resolve confidence values associated with blackboard objects. The KSs are independent specialist computational modules that contain the domain knowledge needed to solve a problem.

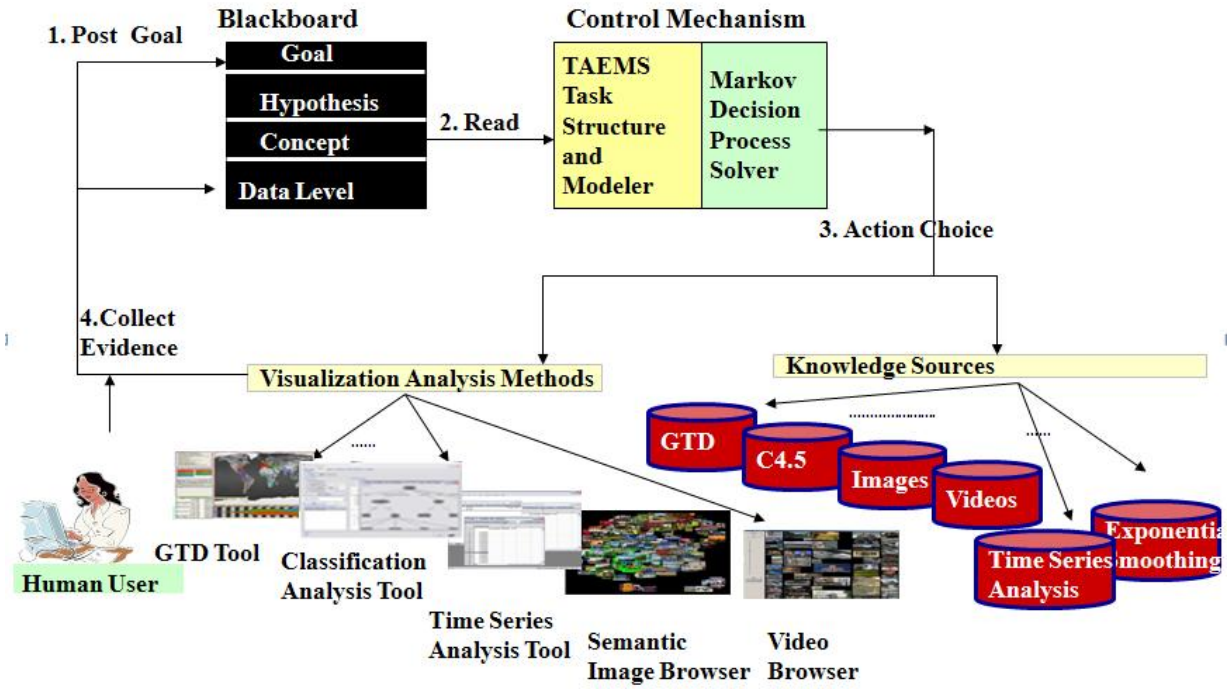

Fig2: RESIN's control flow

As shown in Figure 2, the problem-solving process is initiated

when the Human User posts a goal on the AI Blackboard and this action triggers the agent (Step 1). The goals are domain-dependent and can be split into smaller sub-goals. For example, one sub-goal is to predict the missing information of an input vector containing the partial information about a single terrorist incident. Based on the blackboard's preprocessed result, the TÆM.S. task structure modeler generates an appropriate task structure and translates it to the MDP solver for action assessment (Step 2). Using dynamic programming, the MDP solver computes the optimal policy based on resource constraints (e.g. deadline) and generates the best action, which will trigger appropriate methods to perform predictive analysis (Step 3). Through a built-in user interface, the agent enables the user to interact with the visual analytics tools supporting the mixed-initiative problem solving process, to validate the initial RESIN results and to post their results back to the AI blackboard (Step 4). Using these visualization results as well as previous analysis results, the blackboard will then propagate the evidence information and verify a specific hypothesis with an associated confidence value.

Multi-level Blackboard Database

RESIN's AI blackboard is a global shared repository containing problems, elementary data, a set of partial solutions, contributed information, and other data, which is available to all KSs and serves as a communication medium. It is the kernel of this reasoning agent, providing a reasoning approach for the information that has been discovered and produced. It contains four different levels: Goal, Hypothesis, Concept, and Data, in order of decreasing levels of abstraction. The Goal level stores the goal of the problem and resolution information. The Hypothesis level contains concepts which are represented in the Concept level. The Data level contains the data/evidence gathered to (in) validate the various hypotheses. The layered hierarchy allows for explicit modeling of concurrent top-down and bottom-up processing, while maintaining a clear evidential

path for supporting and contradictory information. The information at a given level is derived from the level(s) below it, and it in turn supports the hypothesis at higher levels.

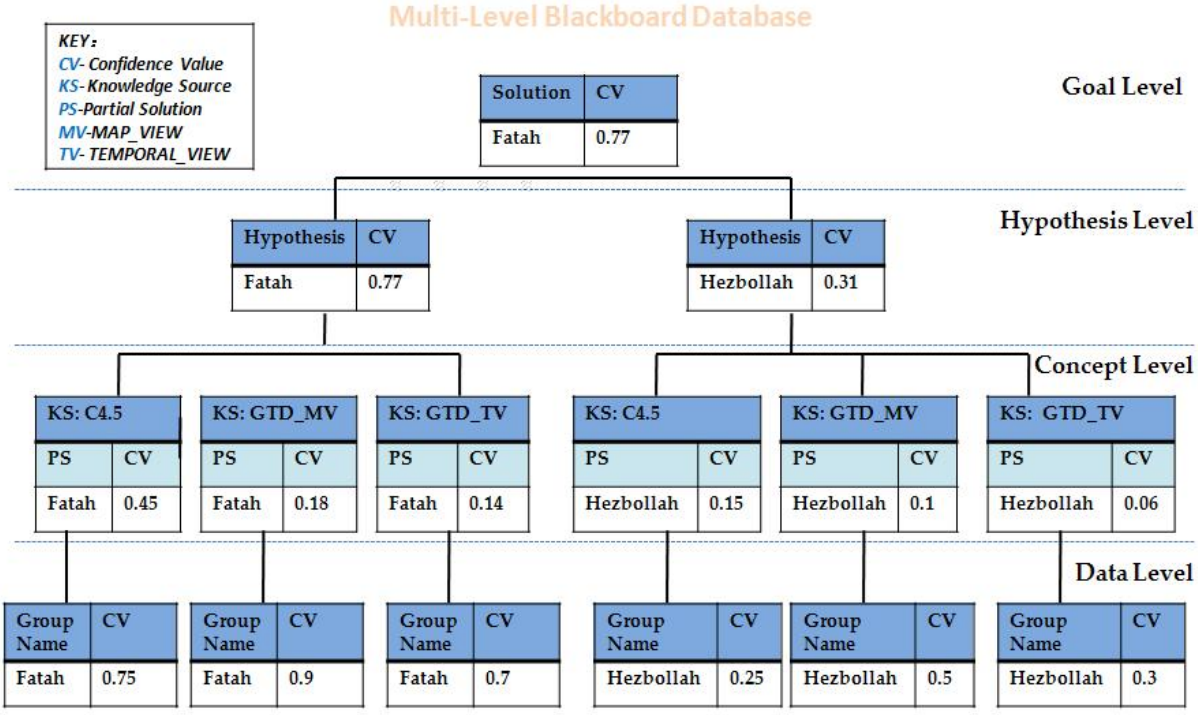

Fig3: Blackboard Data Structure

RESIN in Action

The problem solving process is initiated when the human user posts a goal on RESIN's blackboard, and this action triggers the RESIN agent. The goal is to predict the impact level and count occurrence of incidents perpetrated by a certain group or groups which has perpetrated the current terrorist incident.Given a particular deadline for the problem solving process, RESIN will assist the analyst in performing the predictions for the unknown group name (CP), the occurrence trend for this particular group (EOP), as well as the impact level for future events (EILP).

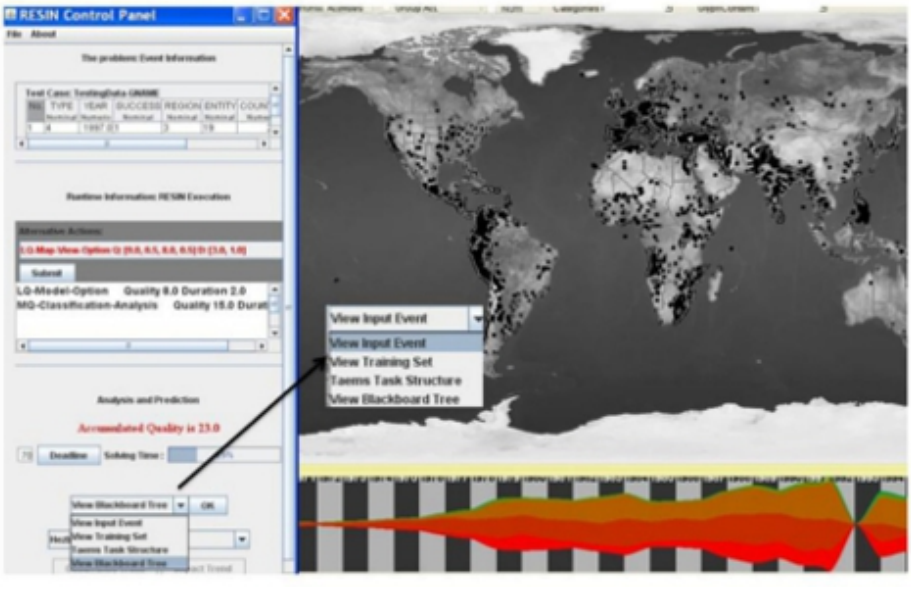

Fig4: RESIN's Control Panel

Fig5: Runtime Views

Empirical Evaluations

This experiment is based on a training set of 2700 incidents selected from the GTD, and for each task we use the same ten incidents from a test set, with different deadlines from 30 to 70 (ranging from a very tight to a loose deadline). Duration describes the amount of time that the action will take to execute. Different applications have different notions of the basic unit of time. The deadline is the hard constraint by which the overall task has to be completed.There are ten users involved in the experiments with access to the GTD tool. Each user will determine the confidence values towards initial predictions, with values from -0.9 (strongly disagree and dispute the result) to 0.9 (strongly agree and accept the result) through interactions with MAP_VIEW and TEMPORAL_VIEW .We compare the predictive performance of the MDP policy and a Deterministic Schedule (DS) for task structures under different deadlines. DS is a deterministic process scheduler that builds a

static schedule with the highest possible quality.

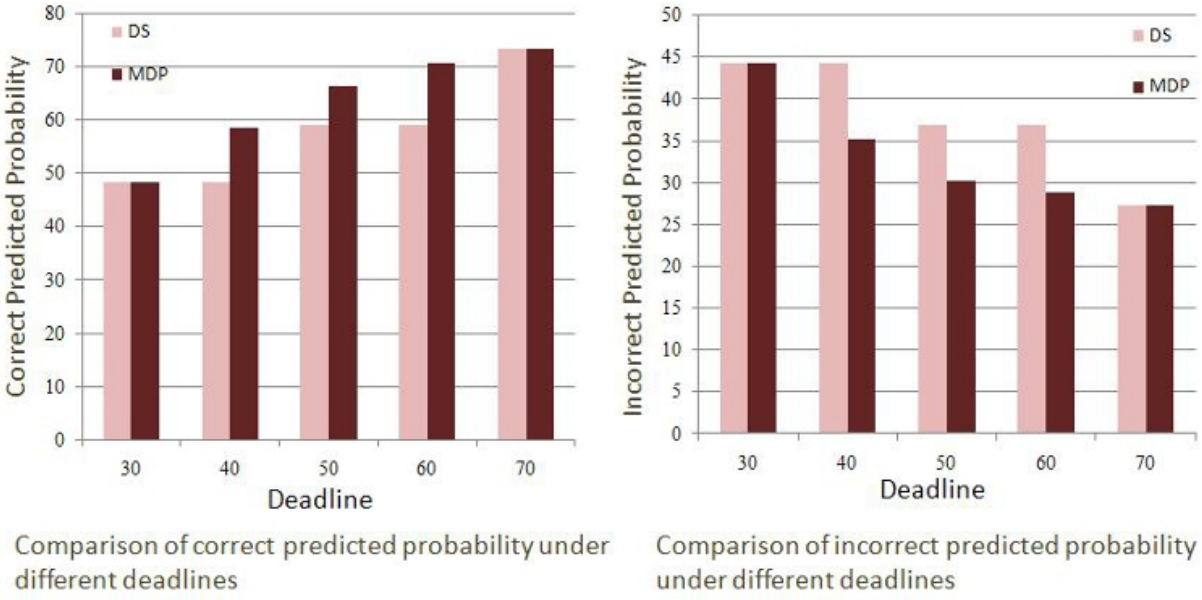

Fig6: RESIN's CP Performance

We also use the monthly data from January 1970 through December 1996 as the training set to create the model on which to predict the occurrences of terrorist attacks from January 1997 through December 1997.

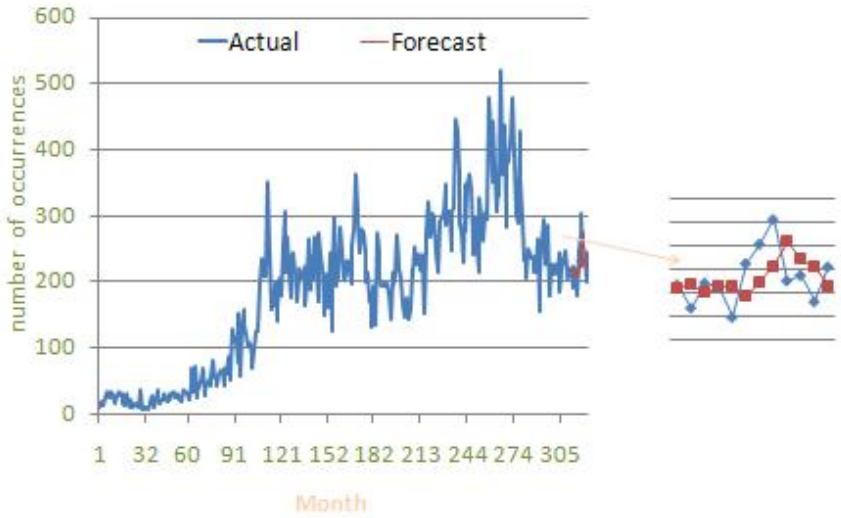

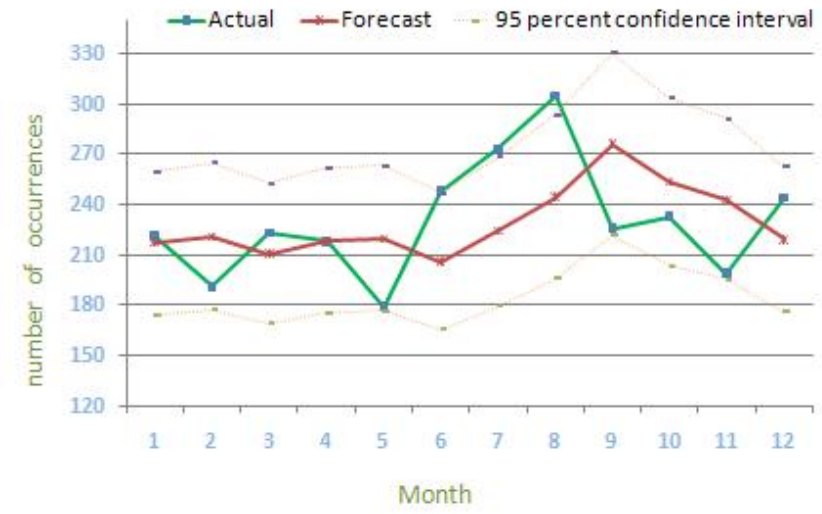

Fig7: Resin's EOP Performance

This work is supported by the grant to establish the Regional Visual Analytics Center (RVAC) at UNC Charlotte and Georgia Tech. Professor Ribarsky's presentation on the vision of the RVAC can be found here.

Links:

- SouthEast Regional Visualization and Analytics Center (SRVAC)
- Charlotte Visual Analytics Center

Publications:

- Jia Yue, Anita Raja and Bill Ribarsky "Predictive Analytics Using a Blackboard-based Reasoning Agent" To appear as Short Paper in Proceedings of IAT-2010.
- Jia Yue, Anita Raja, Dingxiang Liu, Xiaoyu Wang, William Ribarsky " A Blackboard-based Approach towards Predictive Analytics ", Proceedings of AAAI Spring Symposium on Technosocial Predictive Analytics, pp 154-161, Stanford University, CA, March 23-25, 2009.
- Dingxiang Liu, Jia Yue, Xiaoyu Wang, Anita Raja, William

Ribarsky "The Role of Blackboard-based Reasoning and Visual Analytics in RESIN's Predictive Analysis", To appear in Proceedings of 2008 IEEE/ WIC/ ACM International Conference on Intelligent Agent Technology (IAT 2008), Sydney, Dec 9-12, 2008. Extended version is CVC Technical Report CVC-UNC Charlotte-08-29, July 2008 .

- Dingxiang Liu, Anita Raja, and Jayasri Vaidyanath, " TIBOR: A Resource-bounded Information Foraging Agent for Visual Analytics " Proceedings of 2007 IEEE/ WIC/ ACM International Conference on Intelligent Agent Technology (IAT 2007), pp 349-355, Silicon Valley, CA, Nov 2-5, 2007
- Anita Raja and Ashok Goel, "Introspective Self-Explanation in Analytical Agents", Proceedings of AAMAS 2007 Workshop on Metareasoning in Agent-based Systems, pp 76-91, Hawaii, May 2007

.

CONTACT US

Distributed Artificial Intelligence Research Lab
9201 University City Blvd.
Charlotte, NC 28223

Dr. Anita Raja
Office: Woodward 310D
Phone: 704-687-8651
Fax: 704-687-4893

ADDITIONAL CAMPUS RESOURCES