THE COLLABORATIVE LEARNING AGENT (CLA) IN TRIDENT WARRIOR 08 EXERCISE

Charles Zhou, Ying Zhao and Chetan Kotak
Quantum Intelligence, Inc., 3375 Scott Blvd. Suite 100, Santa Clara, CA 95054, U.S.A.

Keywords: Agent learning, Collaboration, Anomaly search, Maritime domain awareness, Intelligence analysis, Unstructured data, Text mining.

Abstract: The Collaborative Learning Agent (CLA) technology is designed to learn patterns from historical Maritime Domain Awareness (MDA) data then use the patterns for identification and validation of anomalies and to determine the reasons behind the anomalies. For example, when a ship is found to be speeding up or slowing down using a traditional sensor-based movement information system such as Automatic Information System (AIS) data, by adding the CLA, one might be able to link the ship or its current position to the contextual patterns in the news, such as an unusual amount of commercial activities; typical weather, terrain and environmental conditions in the region; or areas of interest associated with maritime incidents, casualties, or military exercises. These patterns can help cross-validate warnings and reduce false alarms that come from other sensor-based detections.

1 INTRODUCTION

Port security is important. The Navy needs to enhance its awareness of potential threats in the dynamic environment of Maritime Domain Awareness (MDA) —and plan for potential high-risk events such as use of maritime shipping for malicious activities.

With ever-increasing operations with joint, coalition, non-government, and volunteer organizations require analysis of open-source (uncertain, conflicting, partial, non-official) data. Teams of analysts in MDA may consist of culturally diverse partners, each with transient team members using various organizational structures. These characteristics place increasingly difficult demands on short turn-around, high stakes, crisis driven, intelligence analysis. To respond to these challenges, more powerful information analysis tools can be of great assistance to reduce their workload.

Structured data are typically stored in databases such as Excel or XML files with well-defined labels (meta-data). The unstructured data include free text, word, .pdf, Powerpoint documents, and emails. A large percentage of data remains unstructured despite rapid development of database and data management technologies. Organizations have an opportunity to use unstructured data, if analysis tools can be developed. In the MDA domain, both structured data, e.g. Automatic Information System (AIS) data of monitoring the tracks of vessels, and unstructured data, e.g. intelligence reports from various sources, are important. Anomalies in the structured data such as vessels that are off tracks can be detected using traditional anomaly detection methods. However, it is challenging to analyze the large amount unstructured data that are available. There are a number of extant tools for text mining including advanced search engine (Foltz, 2002; Gerber, 2005), key word analysis and tagging technology (Gerber, 2005), intelligence analysis ontology for cognitive assistants (Tecuci et al., 2007, 2008); however, better tools are needed to achieve advanced information discovery. Furthermore, it is also challenging is to tie the anomalies detected from structured data to the context of unstructured data, which might shed light on social, economic and political reasons for why anomalies occur.

Trident Warrior is an annual Navy FORCEnet Sea Trial exercise to evaluate new technologies that would benefit warfighers. The CLA technology was selected for Trident Warrior 08 (TW08). This paper reports the results from this exercise. In this paper, we report how the CLA technology was applied and evaluated in TW08.

Zhou C., Zhao Y. and Kotak C. (2009).
THE COLLABORATIVE LEARNING AGENT (CLA) IN TRIDENT WARRIOR 08 EXERCISE.
In *Proceedings of the International Conference on Knowledge Discovery and Information Retrieval*, pages 323-328
DOI: 10.5220/0002332903230328

1.1 Agent Learning

Automate human cognitive tasks e.g. detecting and separating anomalous behavior from normal ones, we train synthetic, learning agents to perform tasks like humans. Agent-based software engineering was invented to facilitate information exchange with other programs and thereby solve problems like humans. Multi-agent, distributed networks were developed to provide for an integrated community of heterogeneous software agents, capable of analyzing and categorizing large amounts of information and thus supporting complex decision-making processes. A learning agent defined in this paper is a single computer program, installed in a single computer node, is responsible to learn and extract patterns from data resided locally in the computer and in a specific domain. The agent is dedicated to periodically monitor in the data (structured, unstructured, historical and real-time) and then separate and compare patterns and anomalies (Figure 1).

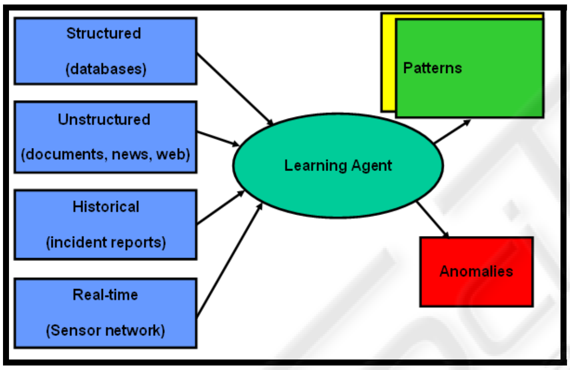

Figure 1: A learning agent ingests structured, unstructured, historical or real-time data and separate patterns and anomalies.

The process is conceptually linked to a full text indexing in the traditional information retrieval. The advantage of the algorithm over the traditional methods is that it captures the cognitive level of understanding of text observations using a few key concepts. Our proposed agent learning algorithm uniquely applies an *anomaly search* method to separate interesting text data from the rest, i.e. separating anomalies and patterns for unstructured data.

Patterns mean something happens more frequently or can be repeated. Anomalies mean something happen less frequently or can not be repeated. As a result of an agent learning process, a learning model is generated to summarize the patterns and anomalies that the agent discovers/learns. Resulting from this process is a learning model containing descriptions of both patterns and anomalies, generated using keywords. Key attributes and statistics are also captured and stored. This process is also referred to as a search index.

1.2 Agent Collaboration

Multiple agents work together to form an agent network. The resulting learning model or index from each individual agent is stored locally in the agent. Each agent can only access and share the learning models or *indexes* of other agents as results of data analysis. However, the original data is not directly shared among agents. A piece of new information is characterized by the collaborative decisions of the patterns or anomalies in all agents in the network.

This is related to distributed knowledge management architecture (Bonifacio, M., et al., 2002). This collaborative infrastructure is a peer-base system, where agent-like applications are distributed among a grid of computers. Each application is considered itself as a peer or node among a network of similar applications. The infrastructure is “fault-tolerate”, “distributed”, and “self-scalable”. With all the advantage of a peer-based system, however, the current peer-based systems lack full-text analysis capability to discover new things.

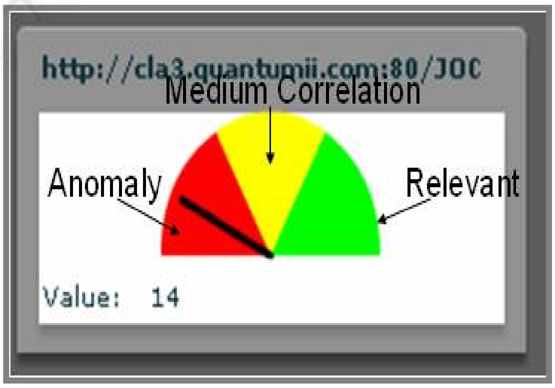

Figure 2: Anomaly Meter.

Agent collaboration is also related to social network research. Social network analysis (Hoff, 2002) is widely used to analyze relational information among interacting units. This framework has many applications in recent years in the social and behavioral sciences including, the behavior of epidemics and dynamics associated with terrorist networks. The social network research is also related to information retrieval and text analysis. For example, the search engine Google uses the interconnectedness of the World Wide Web for page ranking (Brin et al., 1998). Our solution uniquely couples agent learning and collaboration that can significantly increase the automation with desired collective behavior in a decentralized, self-organized environment.

1.3 TW 08 Setup

We used three agents learning patterns from three historical maritime domain information sources. Each agent is responsible for mining information from one collaborative MDA partner such as Navy, Police or Coast Guard as shown in Figure 3. We used open-source unstructured data, i.e. websites, news and freelance reports as the training data.

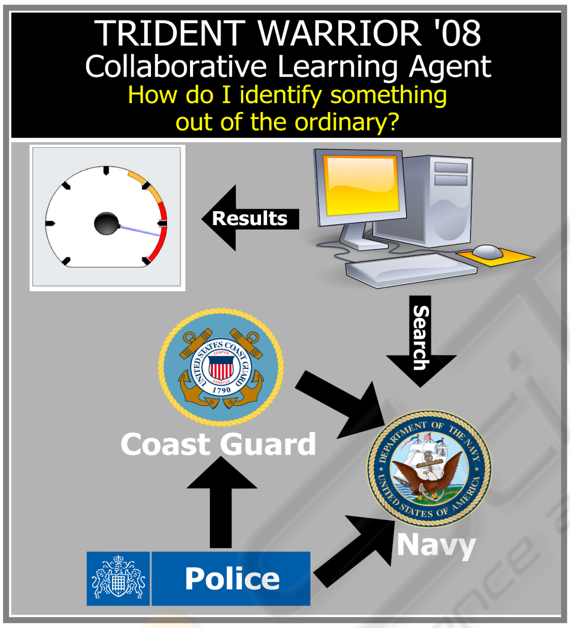

Figure 3: CLA -- Ability to learn from unstructured data and tie the patterns and anomalies with structured data.

We are able to access the Navy real-time vessel AIS data from SPAWAR DS COI (SPAWAR data sharing, community of interest, https://mda.spawar.navy.mil) as shown in Figure 4. The SPAWAR data is in not classified, only requiring a DOD PKI for the access. The MDA DS COI, Automatic Identification System (AIS) track information and associated alerts including data from Navy Organic Sensors aboard Navy ships, The Department of Transportations (DOT), The United States Coast Guard (USCG), Office of Naval Intelligence (ONI) to track merchant shipping. The data is published as the NCES Messaging Service that can be integrated with standard web services. The data shows worldwide real-time ship's names and locations.

Figure 4: real-time AIS data from MDA DS COI.

In a test process, when a piece of real-time information is newly observed, i.e. a ship is observed at a location, it goes through the CLA network; the network then returns a report of anomaly search results which shows if the new information is correlated with the patterns and to what degree the correlation is. In this exercise, an input is each vessel's name and location is identified by AIS is classified into prediction categories (see Figure 1): 1) Anomaly (red), i.e. a search input that has low correlation with previously discovered context patterns; 2) Relevant (green), i.e. an input is highly correlated to the previously discovered knowledge patterns; 3) Medium Correlation (yellow), i.e. between relevant and anomaly; 4) Irrelevant (white), i.e. an input is not related to any of the agents' knowledge patterns, or a correlation value can not be computed from the CLA network.

A user will observe the test process for about 100 real-time inputs. Each input (sequence) represents a vessel's name or real-time location from the SPAWAR MDA DS COI. The input is checked against the patterns in the CLA network to see if anything is of interest or relevance to the vessel or its location; for example, was the vessel seen anywhere else before? Were there any incidents/activities/events reported in the vessel's location? A user will compare samples of the categorizations (i.e. anomaly, relevant, medium correlation or irrelevant) from the CLA network with his/her own categorization.

1.4 Experiment Objective

The objective was to employ a collaborative learning agent (CLA) to derive behavior patterns from historical MDA data, and use patterns in predictive analysis, with context for those predictions. The questions that were needed to be answered related to this objective are listed as follows

- Is the intelligent agent in CLA capable of learning from unstructured, historical information (for example, chat log from all TW participants, samples from NCIS)?
- Is CLA capable of prediction from unstructured data?
- Does CLA predict relevant anomalies or interesting MDA behavior?
- Is CLA accurate when its predictions are compared with predictions from human analysts?
- Are the CLA interface, visualization and display usable?

2 EXPERIMENT DESIGN

This assessment was designed to be made by a CLA Subject Matter Expert (SME) during a period in which the agent "learned" from various sources. Three agents, one for each specific database, were used:

- Agent 1 (http://cla1.quantumii.com/FAIRPLAY) for The Lloyd's Register – Fairplay (LRF) news
- Agent 2 (http://cla2.quantumii.com/JOC) for the Journal of Commerce, which includes information regarding port events, activities, rules, and policies
- Agent 3 (http://cla3.quantumii.com/MPC) for Maritime Press Clippings which are freelance vessel and incident reports.

The CLA analysis process involves three steps:

Step 1: Agent learning: Each agent learns patterns from a single historical data source.

Step 2: Real-time Application: After the learning process, an agent is ready to apply the learning model (e.g. patterns and anomalies) to new data. The agent decides that new data is either anomalous or expected:

Anomalous: An input is an interesting or unique event, for example, a ship or location is not associated with historical location norms.

Expected: An input is a normal or expected event because it fits into the patterns developed by the agent.

Step 3: Agent Collaboration: A set of networked Collaborative Learning Agents (CLAs) forms an agent network and performs a collaboration to decide together if a real-time input is expected or anomalous. Each anomaly is classified into one of four categories using the following rules:

- An input is an Anomaly if all the agents decide the input is anomaly
- An input is Relevant if at least one of the agents decides the input is relevant
- An input is Irrelevant if none of agents decides any relevance
- An input is Medium Correlation if the agents cannot decide if it is an anomaly or relevant.

A collaborative result of the agents is shown in Figure 5, showing critical events are identified – red is an anomaly and green is a pattern.

In order to address the relevant issues involved with each question under the objective, the following approaches were employed for the particular questions.

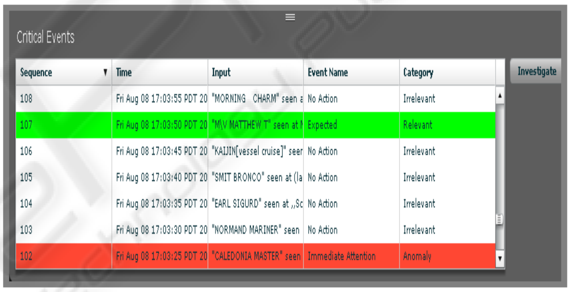

Figure 5: Critical events are identified by agent collaboration. Red event is an anomaly and green event is a pattern.

2.1 Learning

Three agents generated the learning models on June 16, 2008 based on open-source, pre-scenario information through 15 June. A survey was presented to an SME, to address:

Measuring: Are data being ingested from the source?

Measured by: Assessment of data ingested into the training data set, and comparison with sources

Method: Reading the model log, the number of training data points ingested per agent was noted, along with the number of source data points. The percentage of each data set that was ingested was reported.

2.2 Detections and Predictions from Unstructured Data

Three methods were planed for testers or observers as follows:

Method A: Observers answered questions about CLA activity (inputs, gauge activity, additions to critical event tables) that occurred in an 8 minute period. Observations were made to indicate that the dashboard was receiving data, communicating with other agents, and analyzing information.

Method B: Correlation values were recorded in each of 3 agents after an 8 minute time period. Observers documented that agents collaborated in real-time to make decisions and classifications of inputs.

Method C: Observers assessed the results from CLA against his/her own domain knowledge. Questions address usefulness and relevance of the data and whether or not it's "out of the ordinary," i.e., unexpected.

2.3 Anomaly Prediction Relevance

Anomaly prediction relevance was based on the assessment of the observers.

2.4 Comparative accuracy

Accuracy was defined as percentage of correct vs. false positive and false negatives following a post-scenario validation.

2.5 Usability

Usability was defined to be the analysts' assessments of

- Clarity of display
- Extent to which trusted
- Ease of accessing the detailed data.

2.6 Data Collected

Electronic data and observer questionnaires were the basis for evaluation of this approach.

3 EXPERIMENT RESULTS

3.1 Learning

The percentage of training data from individual sources ingested to CLA ranged from 60% to 78%. Some of the data was automatically pruned away because it did not contain relevant contextual information.

The three agents used appeared to learn from the databases and were able to develop patterns within the data.

The consistency of these patterns compared to those that an expert might develop over time was not assessed, but would be possible in future demonstrations.

3.2 Detections and Predictions from Unstructured Data

Observers answered the designed questions as follows:

Method A question: Do you see ship names and/or locations in the Input column in the Critical Events Table? 4 out of 4 (100%) answered yes. All observers noted dashboard reactions and gauge changes indicating that the system was receiving real-time data feeds. The critical event table data was updated during operations, indicating some degree of the detection of anomalies or expected events.

Method B questions: Have you noticed the agent gauges move? 3 out of 4 (75%) answered yes. Have you noticed data being added to the critical event table? 4 out of 4 (100%) answered yes. Observers noted that the correlation values changed during real-time operation, indicating possible collaboration between agents while classifying inputs or developing decisions.

Method C's data (expert assessment of relevance) was used in developing the relevance of anomaly predictions and comparative accuracy of predictions.

3.3 Anomaly Prediction Relevance

These values were analyzed by comparing the ratings of items (relevant or not relevant) by the CLA with those of SMEs.

The CLA identified 44% of the total number of relevant items consistent with experts.

The CLA identified 71% of the total number of non-relevant items consistent with experts.

3.4 Comparative Accuracy

These values were analyzed by comparing the ratings of patterns (high or low correlation with known patterns) by the CLA with similar ratings by SMEs.

The overall accuracy for the CLA predictions was 72%. The overall error rate was 36%. The false positive rate was 53%. The false negative rate was 23%.

3.5 Usability

Usability was determined using surveys to assess the subjective opinions of users. Opinions were generally neutral but divided about the usability of the CLA system. This is not unexpected, as this technical capability was completely new to users, and work will have to continue in order to integrate and implement this category of capability.

4 CONCLUSIONS

Considering the problem of MDA a challenging and highly complex environment, CLA achieved unique results in automating learning from the immense but relevant information that emerges from the unstructured environment which continually refreshes the information domain with new and unstructured data. CLA used the agent technology in new ways, adds to “sense-making” capabilities of the future.

ACKNOWLEDGEMENTS

This work was supported the Navy contract N00244-08-P-2638.

REFERENCES

Bonifacio, M., et al., 2002. A peer-to-peer architecture for distributed knowledge management. in the 3rd International Symposium on Multi-Agent Systems, Large Complex Systems, and E-Businesses.

Brin, S. and Page, L.,1998. The Anatomy of a Large-scale Hypertextual Web Search engine.” In Proceedings of the Seventh International World Wide Web Conference.

Foltz, P.W., 2002. Quantitative Cognitive Models of Text and Discourse Processing. In The Handbook of Discourse Processes. Mahwah, NJ: Lawrence Erlbaum Publishing,

Gerber, C., 2005. Smart Searching, New technology is helping defense intelligence analysts sort through huge volumes of data. In Military Information Technology, 9(9).

Hoff, P.D., Raftery, A.E. and Handcock, M.S., 2002. Latent Space Approaches to Social Network Analysis. Journal of the American Statistical Association, 97. http://elvis.slis.indiana.edu/fetched/article/1129.htm

Tecuci G., Boicu M., 2008. A Guide for Ontology Development with Disciple, Research Report 3, Learning Agents Center, George Mason University.

Tecuci G., Boicu M., Marcu D., Boicu C., Barbulescu M., Ayers C., Cammons D., 2007. Cognitive Assistants for Analysts, in John Auger, William Wimbish (eds.), Proteus Futures Digest: A Compilation of Selected Works Derived from the 2006 Proteus Workshop, Joint publication of the National Intelligence University, Office of the Director of National Intelligence, and US Army War College Center for Strategic Leadership.