2008 CCRTS

C2 for Complex Endeavors

A Cognitive-based Agent Architecture for Autonomous Situation Analysis

Topic(s):
Cognitive and Social Issues
C2 Architectures

Gary Berg-Cross
EM&I
455 Spring Park Place 350
Herndon, VA, USA
Gary.berg-cross@em-i.com

Augustine J. Kwon (Principal point of contact)
ICF International
3100 Presidential Drive 300
Fairborn, OH, USA
937-429-8660
akwon@icfi.com

Wai-Tat Fu
University of Illinois at Urbana-Champaign
405 N. Mathews Ave
Urbana, IL, USA
wfu@uiuc.edu

ABSTRACT

A Cognitive-based Agent Architecture for Autonomous Situation Analysis

This paper presents a preliminary investigation on a cognitive basis for large scale, distributed intelligent agents with some level of autonomy. We believe that the proposed architecture can better address many known problems of machine based situation understanding. We propose a practical and economically feasible solution of intelligent agents adapting existing agent modeling frameworks, ontologies from semantic web technology as well as reasonable situation domain models. These can be brought together with a suitable cognitive architecture ACT-R (Adaptive Control of Thought—Rational) which could be used to provide key roles in more human-like situational awareness capability in emergency and disaster operations, especially where sensor information is harvested from semantically heterogeneous data sources. Existing situational ontologies and vocabularies can be supplemented by using DOLCE (Descriptive Ontology for Linguistic and Cognitive Engineering)'s formal ontology. This serves as a metalevel ontology that can relate different ontology modules and can generate new categories to extend ontology using agent learning as needed. Semantically-rich Descriptions & Situations ID&S) ontology can also be utilized to provide a theory of ontological contexts that can describe various types of context including non- physical situations, plans, and beliefs, as entities so they can be effectively communicated and understood among agents in order to reach a common situational view. Lastly the paper discusses the potential performance measurement issues and future research for this particular solution.

1. Introduction

As witnessed by the diversity of research papers in past artificial intelligence (AI) related conferences, no single AI technique or tool has been found to be adequate to address all the characteristics presented by a human situation awareness process. This is true even for relatively simple software based information agents using the meaning of the data for information sharing such as envisioned in the original Semantic Web (SW) concept. Currently, there are multiple views on a suitable architecture and the nature of the knowledge needed to develop successful SW agents. Moreover, agent-based information integration, as typically discussed for the SW is not the only type of information fusion being actively researched. Sensors provide systems access to real-time (or near real-time) streams of actual, low-level data of events. These serve as input to other agents which both collectively or individually fuse these sensor data and integrate them with higher level situational knowledge. Such sensor-based processing is used in many domains, including disaster response, crisis management, modern battlefield operations, and health monitoring. These situations are characterized by multiple, distributed heterogeneous information sources, and rapidly changing situations that may include mobile agents/objects. Special agent capabilities are needed because situations involve a large number of inter-dependent, dynamic objects that change their states in time and space, and engage each other in fairly complex relations. As a result, required intelligent systems capabilities include effective methods for situation recognition, prediction, and reasoning activities. Collectively, these capabilities have been called situation management [1].

A task to build such situation understanding agents requires processing dynamic situations using complex cognitive modeling, design and population of formal situation ontologies, collection, and fusion of sensor data. Current agent systems still have difficulty accommodating things like diverse spatiotemporal information within a single analytic context in a suitable period of time. Yet, as part of analytic process for understanding situation, humans easily integrate both quantitative and qualitative information assessments to quickly arrive at analytic conclusions. The discrepancy may, in part, be in part due to the duality behind human cognitive architecture. The dual processing theory [2] distinguishes between cognitive processes that are:

- Unconscious - rapid, automatic and high capacity
- Conscious - slow and deliberative

This characterizes human reasoning as a robust interplay between an easily believed perception-based system, and a more cognitively demanding logic-based reasoning system. This is a view of intelligence as a developed phenomenon that balances multiple reasoning mechanisms, together with scruffy modules of knowledge which learn to deal with situations that are only partial predictability, due to dynamics and the absence of precisely defined states. Berg-Cross [3] [4] suggested that a multi-level hybrid architecture—based on a cognitively realistic foundation—could approximate human performance for diverse analytic problems. To be practical, such architecture would build on the existing agent models, semantic web technology and standards, as well as reasonably adequate knowledge and domain models. The direction proposed herein is an architecture that leverages recent advances in machine learning, distributed agent technology, and rich semantic representation brought together within a suitable

cognitive architecture—ACT-R (Adaptive Control of Thought—Rational). ACT-R is particularly suitable for performance measurements of intelligent capability because it has been widely shown to be capable of simulating human performance on a range of task and matching human learning improvement profiles over time.

A part of paper discusses why ACT-R was chosen over other cognitive architectures and how ACT-R architecture can be used to learn and perform aspects of situation assessment. Our discussion of work is divided into three parts. First, we describe what is understood about situation understanding cognitively, and the ingredients of a situational ontology are described. Simple domain ontologies can be enhanced by leveraging some foundational ontology and their modules to formalize concepts like 'participation'. Framing models like Description of Situations (D&S) can be used to model how knowledge of information can be shared by agents. Second, the ACT-R architecture is described showing how it might be populated and trained on situational understanding tasks. A third and final section summarizes the feasibility of the approach and describes future research and development for the proposed solution. The final section also addresses foreseeable performance measurement issues with the proposed solution.

2. Situational Knowledge and Ontologies

As intelligent agents anchored in the real world, we generally understand the idea of 'situations' and can be said to have 'situational awareness' (SAW). Automation of SAW using agent architectures has proven increasingly important for complex designs as diverse as air traffic control, power plant operations, and disaster management [6]. A rational empirical approach to SAW & understanding is generally defined with three sequential components: (1) perception/awareness of elements/objects in the environment within a volume of time and space; (2) along with a comprehension of their functional nature and organizational relationships (their 'meaning'); as well as (3) an ability to go beyond SAW to project the status & relations of situated objects in the near term as an empirical test of 'expectations'. A top-down, rational model of SAW incorporates an agent's goals & objectives into its reasoning about events, relations, and situations. This helps upper-level agents reduce the number of possible relations definable within an agent's knowledge to constrain situational possibilities. By knowing something about what is expected, attention on relevant events and relations can improve agent operation [5].

Situation/context-aware systems have been proposed as an important class of applications and an important step towards ubiquitous computing. Examples of such agent systems described at the 1st International Workshop on Agent Technology for Disaster Management [6] include discussion of agent architectures to handle coordination via intents, multi-agent learning to support urban planning, and training based on agent-based situation simulation. Often, such work involves sensor-based data being fused into situational information. A typical sensor based multi-agent architecture consists of several functional layers: a set of dedicated sensor agents are usually responsible for the initial degree of sensor data processing and they are considered the first layer; the next set of agents are responsible for associating output from these sensor agents with object concepts derived from domain knowledge and inference; another set of 'higher' agents are responsible for transforming associated information into 'situations.' Such systems assemble operating pictures using precision geospatial environment information layers

(modifiable digital overlays) that can support decision making based on the detailed 'knowledge' shared by the agents sensing a physical environment. An example of a such a system called SAPPHIRE (Situation Awareness and Preparedness for Public Health Incidents using Reasoning Engines) shown below dealing with Public Health Incidents [7]. Pollution sensing includes toxic gases such as carbon monoxide (CO), sulfur dioxide (SO_2), hydrogen sulfide (H_2S), and nitric oxide (NO), etc., with a half dozen meteorological factors and chromatography data such as ethane and ethylene, etc. Systems like SAPPHIRE are not limited to direct sensor feeds, and can include entire reports of other agent's processing (hence NLP capability as shown in the Figure 1). In such architectures, agents at higher levels in the "network" may have more

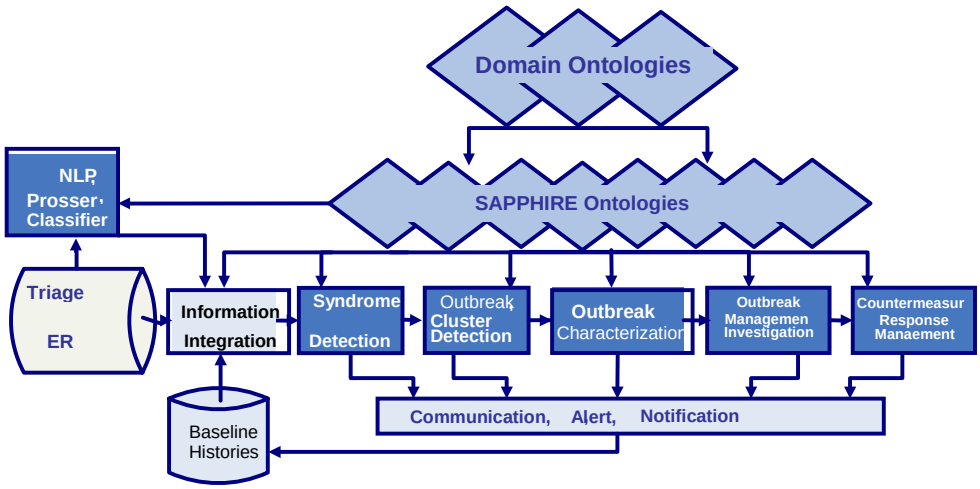

NLP – Natural Language Processing
ER - Emergency Room

Fig. 1: SAPPHIRE [7]

responsibility whereby they may use more elaborate communication protocols, taking into account and monitoring the information provided by lower-level sensor agents. Several things are needed to make such systems effective. For sensor data fusion, common sensor standards are needed to create a common sensor data model. To be able to handle conflicting information from sensors, higher agents might need a capability for incremental, flexible perceptual/ conceptual learning, which is feasible through retrieval of relevant memories.

Generally multi-agent architectures include adaptive infrastructures that can be used to make problem solving responsive if not adaptive to some degree real-time constraints, available network resources, and coordination requirements. Multi-agent architectures often include specialized components for adaptive functions such as local agent scheduling, multi-agent coordination, organizational design, detection and diagnosis, and on-line learning, a core characteristic of adaptive architectures. These can be configured to interact so that a range of different situation-specific coordination strategies can be implemented and adapted as problem 'situations' evolves. A common model to support these is that situations are represented as a high-order knowledge type of concept that is formed using existing concepts. We assume that situation knowledge is formed by an agent's interaction history with the environment and that agents can form situation "concepts" by "observing" that certain patterns of sequence of inputs from the environment. Thus for agent adaptive learning evaluative feedback should follow a

certain action/class of actions, or the next input event from the environment given the current action. For all these reasons robust knowledge is needed to support adaptive agent reasoning.

Knowledge structured as ontologies are increasingly used to define vocabularies for the Semantic Web (SW) thrust to allow intelligent agents to exchange queries and assertions among agents [8]. More recently, ontologies have been applied to represent the more complex problem of sensor-driven situational understanding [9]. We illustrate the use of two levels of ontologies for understanding and representing situational knowledge – a mid-model of situations and a more foundational model that captures more of the event aspects of participation in situations as well as the relationship between descriptions and situations as meta-knowledge which can be used by agents. The SAW 'model' [5] is a 'light' ontology capturing the core elements and relations of a rational agents view of situations, as shown in Figure 2. This provides a principled way to describe situations including crises and disasters.

In the SAW ontology model, there are primary classes: SituationObject, PhysicalObject, and Events. The organizing point in the ontology and resulting models is defined as a relationship to three things: Goals, SituationObjects, and Relations. SituationObjects are entities in a

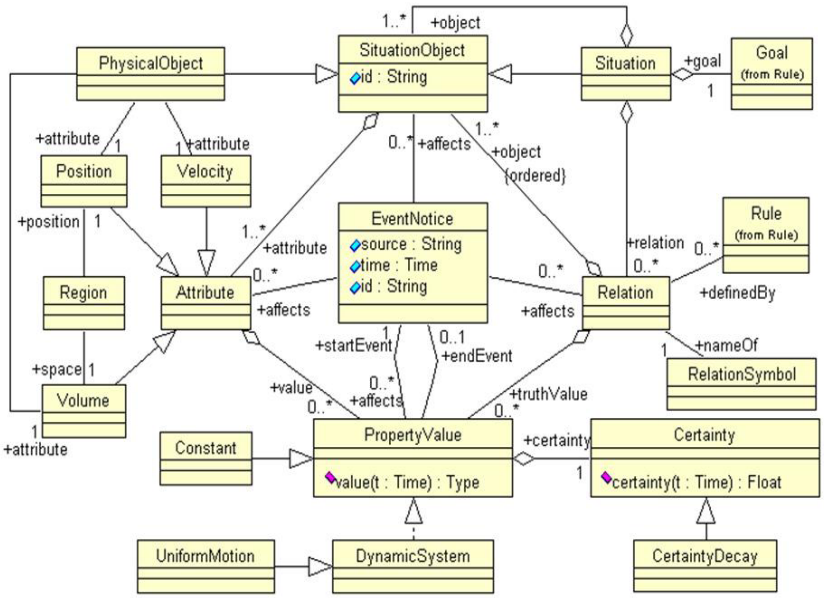

Fig. 2: Core Concepts in the SAW Ontology [5]

situation that participate in relations and can have characteristics (i.e., Attributes). These attributes define values of specific object characteristics, such as expected/unexpected, weight or color. Objects may be PhysicalObject (a sub-type) with Volume, Position, and Velocity. Goals Relations are used in a structuring geometrical (not well represented in SAW) and positional aspects of these concepts. For example, to represent bus transportation systems for flood situations being reported on by environmental sensors, we need both realistic street information that can be overlaid by bus stops and also more abstract timetables along with timetables and circumstances that apply in emergencies as sensors report water height, closed streets, etc. This would allow us to represent situations about stops on a bus route that is 'near' the intersection of two streets. However, in the event of a flood this information will be supplemented by

elevation information and perhaps closeness to streams. Such features are not typically captured in geospatial databases and processable by Geographic Information System (GIS) functions. In our architecture, this information would be integrated by a situational agent using components of information from sensor agents and also making use of geospatial repositories of information. To handle these requirements the SAW ontology can be improved by grounding it in a more foundational ontology like DOLCE (Descriptive Ontology for Linguistic and

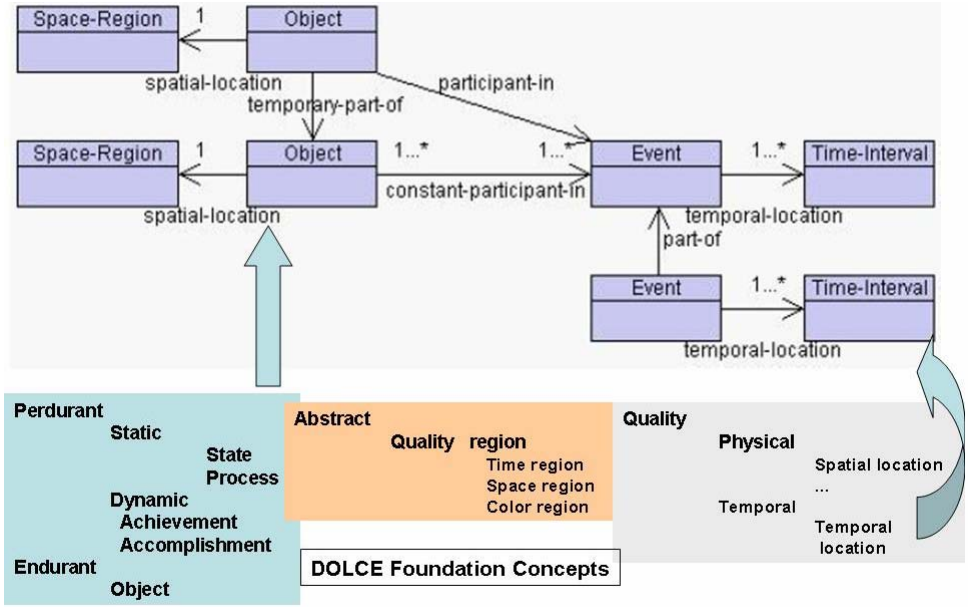

Fig. 3: Ontology Design Pattern for Object 'Participation' in Situations Built on DOLCE Foundation

Cognitive Engineering) developed within the WonderWeb Project (European Union Funded Project 5) [10]. DOLCE is a cognitively based, 'reference' ontology, consisting of about 30 classes, 80 properties and many more axioms.

It was designed to provide a sufficiently neutral base to map, integrate, and build domain ontologies, such as an improved SAW ontology. DOLCE includes the idea a high-level participation pattern of objects taking part in the Events on the SAW model. DOLCE conceptualizes endurants (Objects or Substances) and perdurants (Events, States, or Processes) as distinct types linked by the relation of 'participation'. Participation patterns help us understand the structure of repeating events that occur for types of situation. As shown time in Figure 3, time in indexing is provided by the temporal location of the event at time *interval* duration, while the respective spatial location at a *space region* is provided by the participating object. The general pattern in Figure 3 uses an extended version of UML which can be converted to the DOLCE-time-plus [11] light ontology and is intuitively applicable to situations of interest to us including disasters, and health monitoring. Particularly nice for representing knowledge that agents need to fuse information into situations is DOLCE's use of an Information-object design schema/pattern as shown in Figure 3. This model formalizes how descriptions, which one agent may receive from another, serve as Descriptions for Situations. This extension to the DOLCE foundational ontology is called the Descriptions and Situations

(D&S) Ontology and can be used to define agent workflow - how the information in messages can be used by an actor who play a specific role.

The D&S ontology [12] is based on a first-order logic conceptualization that supports manipulation of *descriptive objects* (such as clinical plans, evacuation routes, emergency plans, institutions, etc.). This allows representation of hypotheses about *situations* (such as cases, facts, settings)[1]. D&S's explicitly committed conceptualization as a distinction between an *unstructured world* or *context*, and *intentionality* (an agent description) that recognizes a *structure* (situation) in that world or context. One nice thing about ontological commitment using D&S is that it supports organizing domain theories for areas like disasters & healthcare into different ontologies as well as into different *descriptions* or *situations*. For example, "a flood situation" is a disaster entity whose conceptualization is described and realized in several modules—disaster, transportation, hydrology, geography etc. Finally, the DOLCE models can serve as modular, meta-level ontologies that can relate different ontology modules and can generate new categories to extend ontology (by agent learning) as needed.

3. Rationale for ACT-R

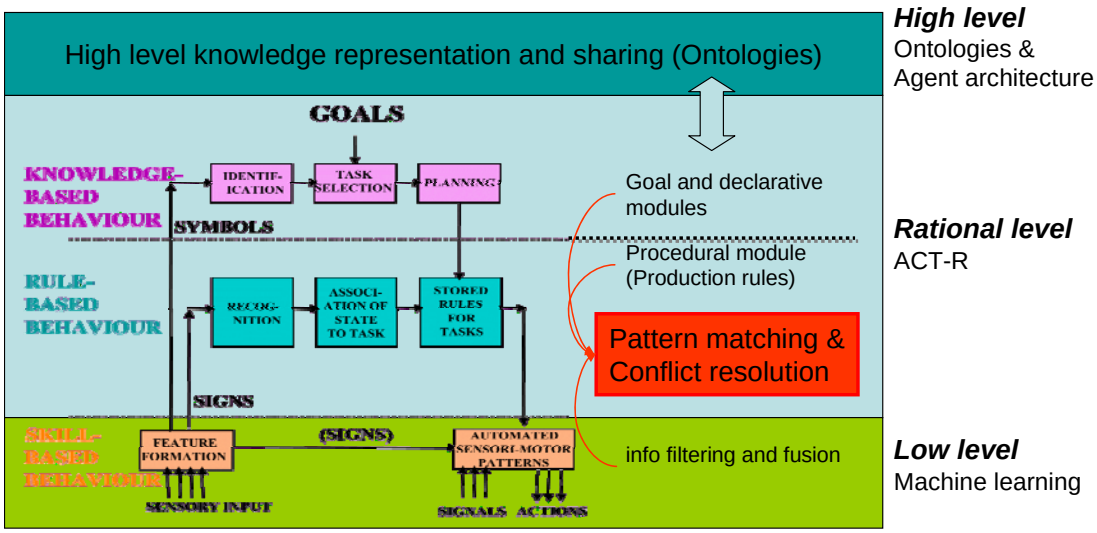

Fig. 4: Adapting ACT-R to a Multi-level Hybrid Approach

Two major production-system based cognitive architectures, ACT-R and SOAR [13], were considered to provide much needed rational level functionality for the proposed situational awareness architecture. Ultimately the team decided to go with ACT-R due to its closer resemblance to human cognition process. The major difference between ACT-R and SOAR is that ACT-R has been constantly evolving along the direction that makes it closer to the human-level intelligence, in the sense that it closely resembles the perceptual and cognitive

[1] When D&S plugged into DOLCE it results in 'DOLCE+' with description being a non-physical endurant. A situation is added as a top level.

constraints by humans. Achieving human-level intelligence is always more challenging than developing powerful algorithms, as years of research in artificial intelligence show that while we can build *deep blue* to beat the world's best chess player, we still cannot develop an artificial systems that will play chess *like* the world's best chess player. In fact, recently rigorous tests on the ACT-R architecture has been done through advanced techniques in brain imaging, providing solid evidence that the architectural constraints closely match the neuro-physiological mechanisms of humans. We consider these constraints as strong theoretical claims that increase the likelihood that *correct* models of human-level intelligence can be constructed, at the same time reducing the likelihood that powerful but *incorrect* models will be developed. This is a very important rationale for our decision on using ACT-R: Although one may argue that SOAR may be more "theoretically powerful" by covering psychologically unreal possibilities, we consider that a weakness, not strength. As a close match to human-level intelligence is critical in ensuring the technologies developed will match well with the *real* perceptual and cognitive profiles of humans who are interacting with these technologies.

Under the proposed multi-agent based framework, the primary functionality of individual agents can be determined by one or more plugged-in ACT-R models. This can use ontology translators similar to those proposed by Wray et al. [14] to allow various ACT-R models and their host agents to share domain knowledge, along with inference knowledge captured in various ontologies. A system of different intelligent agents involved in various steps of the SAW process can work collectively to achieve a common system goal. Each ACT-R model has its own set of knowledge representation, input/output modules that allow it to interact with the external world as well as with other agents/models, and has its own learning mechanisms that allow it to adapt to the new situations and environmental conditions. Since ACT-R is developed based on the principle that knowledge are always rationally deployed to decide on the next set of actions, each model can behave rationally based on the existing knowledge it has. For example, ACT-R has a built in Bayesian learning mechanism that allows it to retrieve the relevant knowledge structure based on its need probabilities at a particular context of situations [15], as well as a reinforcement-like learning mechanism that learns to adapt to the statistical structures of the environment so that actions that have led to successful outcomes before will more likely be selected in the future [16]. These sophisticated mechanisms allow each ACT-R model to gradually adapt and learn the skills and knowledge required for different situations.

Another important advantage of the ACT-R architecture is its ability to span components of cognition that have traditionally been treated as separate in cognitive psychology. For instance, a model of Intelligence, Surveillance, and Reconnaissance (ISR) analysis will likely involve reading and language processing, spatial processing, memory, problem solving, reasoning, and skill execution and acquisition. ACT-R not only has a generic knowledge representation across these different components of cognition, but it also specifies how these components are integrated to produce behavior. At its lowest level, ACT-R has both a spreading-activation mechanism to predict accessibility of declarative knowledge and a reinforcement-learning mechanism to predict the future success of certain actions.

For example, knowledge that is needed frequently or repetitive actions with certain outcomes will eventually lead to skilled behavior that can be deployed with little cognitive resources. At the middle level, deliberate acts such as attending to relevant parts of the environment or

pressing the right key on a device requires intelligent integration of multiple sources of information and background knowledge to generate intelligent behavior. To this end, ACT-R has a set of 'sub-symbolic' mechanisms that arbitrate how various sources of information should be integrated, and how and what actions should be selected and executed at different situations. Mechanisms at this level are found to be critical to various training goals. For example, work on computer-generated forces [17] shows that training for people interacting with and against synthetic partners is effective only when these agents perform elementary actions like real people. Similar results were obtained by Jones, et al. [18] and their study shows that training is effective only when synthetic pilots make turns with the timing of real pilots.

A cognitive agent is therefore essential to ensure that the simulated environments appear 'real' during training of operators. At the highest level, long-term knowledge are stored as a large set of declarative memory elements and procedural rules. This set of knowledge can be obtained through a diverse set of training scenarios in various situation analysis environments. Taken together the direction proposed here is a multi-level approach that leverages our understanding of cognitive agent architecture in integrating three levels of information processing behavior as shown in Figure 4. In the high level a distributed agent architecture such as Cougaar (cognitive agent architecture) [19] and foundational ontology such as DOLCE will likely work together to provide a means to process high level situational knowledge that may requires immediate attentions. At the rational level, ACT-R will act as a bridge to connect high level and low level situational information processing behaviors through its proven strength in pattern matching and conflict resolution. At the low level, an unconscious behavior such as machine learning will help to transform sensor fed raw data to ongoing situational knowledge through data filtering and fusion. Under the proposed hybrid approach, the anticipated major contribution of ACT-R will likely come from a rational level where most rule-based behaviors through a human-like cognitive capability take place. One of our research goals is to capitalize on the success of ACT-R in simulating the rational/adaptive nature of human information processing to coordinate activities in low level information fusion/selection and high level semantic ontological reasoning to support distributed decision making process in autonomous situation analysis.

4. Conclusion

The first level of performance analysis helps to understand cognitive criteria underlying success with SAW and pointed out potentially problematic areas and real-time issues with agent knowledge which can be addressed by improved ontology. Performance measurement of a cognitive based distributed multi-agent system (MAS) offers unique challenges that must be addressed explicitly in its agent infrastructure. Our study identified several key issues with current agent technology that must be addressed in the future study. These are:

- Definition and selection of appropriate agent social behavior (competitive vs. collaborative) based on current operational conditions;
- Knowledge storages and retrieval – distributed vs. replicated;
- Knowledge sharing and transfer through proper ontologies;
- If knowledge is replicated locally, how to deal with synchronization issues if the bandwidth is limited;
- Suitable consensus building mechanism among agents under a competitive mode;
- Hierarchical vs. Peer-to-peer infrastructure or mix – dynamic goal switching for ACT-R;

- Autonomous task delegation and volunteer capability;
- Appropriate performance metrics to evaluate appropriate collective intelligent behavior including social behavior of multiple intelligent agents.

A study done by Helsinger, et al. [20], demonstrated the feasibility of using an performance metrics within Cougaar to improve an individual agent's performance and adaptability to their dynamic operational environment. The challenge and future research for our proposed solution lies in how to overcome these issues of current agent technology so that collaborative distributed situation analysis can no longer be viewed as a computationally challenging and inefficient alternative. One possible key to success is how to maximize the use of cognitive capability in such system, using ACT-R like cognitive plug-ins models. The team believes that further development and refinement of proposed approach will eventually lead to a more human-like and robust autonomous situation analysis platform that can effectively operate in a highly mobile and distributed (both computationally and functionally) environment with very little human supervision.

References

[1] G.. Jakobson, B. Gabriel, J. Buford, and L. Lewis, "A Framework of Cognitive Situation Modeling and Recognition", Military Communications Conference, 2006. MILCOM 2006, Washington, DC, Altusys Corp., Princeton, NJ, Oct. 2006

[2] Jonathan Evans, *Dual-Processing Accounts Of Reasoning, Judgment And Social Cognition* in Annual Review of Psychology (2008, in press)

[3] Gary Berg-Cross, "A Pragmatic Approach to Discussing Intelligence in Systems", Performance Metrics for Intelligent Systems (PerMIS) conference 2004.

[4] Gary Berg-Cross, "Developing Knowledge for Intelligent Agents: Exploring Parallels in Ontological Analysis and Epigenetic Robotics", (invited paper) Performance Metrics for Intelligent Systems (PerMIS) conference 2006

[5] C. Matheus, M. Kokar, K. Baclawski, J. Letkowski, C. Call, M. Hinman, J. Salerno, and D. Boulware, "SAWA: An Assistant for Higher-Level Fusion and Situation Awareness", In Proc. SPIE Conference on Multisensor, Multisource Information Fusion, pp. 75-85, 2005

[6] Nicholas R. Jennings, Milind Tambe, Toru Ishida, and Sarvapali D. Ramchurn, First International Workshop on Agent Technology for Disaster Management, Hakodate, Japan 8th May 2006

[7] Parsa Mirhaji, MD & Robert Coyne, "The Semantic Web and Health Information Systems", SICoP Conference 2, 25th April 2007

[8] Heflin, J., Hendler, J., and Luke, *S.*, *SHOE: A Blueprint for the Semantic Web*, Fensel, D., Hendler, J., Lieberman, H., and Wahlster, W. (Eds.), Spinning the Semantic Web. MIT Press, Cambridge, MA, 2003

[9] M. Kokar, C. Matheus, K. Baclawski, J. Letkowski, M. Hinman, and J. Salerno, "Use Cases for Ontologies in Information Fusion**",** In Proceedings of FUSION'04, Stockholm, Sweden, pages 415-422, June 2004.

[10] EU FP5 (European Union Funded Project 5) WonderWeb project (http://wonderweb.semanticweb.org) by the Laboratory for Applied Ontology (http://www.loa-cnr.it.)

[11] Aldo Gangemi, Stefano Borgo, Carola Catenacci, and Jos Lehmann, *Task taxonomies for*

knowledge content, Metokis deliverable d07, June 2004.
[12] Gangemi, A. and Mika, P. "Understanding the semantic web through descriptions and situations", DOA/CoopIS/ODBASE 2003 Confederated International Conferences DOA, CoopIS and ODBASE, Proceedings, LNCS. Springer, 2003
[13] Laird, J., A. Newell, and P. Rosenbloom, "SOAR: an Architecture for General Intelligence", Artificial Intelligence, Vol. 33, Issue 1, pp. 1 – 64, 1987
[14] Wray, R. E., Lisse, S., and Beard J., "Ontology infrastructure for execution-oriented autonomous agents", Knowledge Engineering and Ontologies for Autonomous Systems 2004 AAAI Spring Symposium, Vol. 49, Issues 1-2, pp. 113-122, 30th November 2004.
[15] Anderson, J. R. & Lebiere, C., *The atomic components of thought*, Mahwah, NJ: Erlbaum, 1998.
[16] Fu, W-T. and Anderson, J. R. "From recurrent choice to skill learning: A reinforcement-learning model", Journal of Experimental Psychology: General, vol. 135, no. 2, pp. 184-206, 2006.
[17] Pew, R. W. and Mavor, A. S., *Modeling human and organizational behavior: application to military simulations*. Washington, DC: National Academy Press, 1998.
[18] Randolph M. Jones, John E. Laird, Paul E. Nielsen, Karen J. Coulter, Patrick G. Kenny, and Frank V. Koss, "Automated Intelligent Pilots for Combat Flight Simulation", AI Magazine Vol. 20, No. 1, Pages 27-41, 1999
[19] Cougaar (Cognitive Agent Architecture), a java based open-source agent architecture funded by DARPA (1996) for construction of highly scalable distributed agent-based applications, http:// www.cougaar.org/
[20] Helsinger, A., R. Lazarus, W. Wright, and J. Zinky, "Tools and techniques for performance measurement of large distributed multi-agent systems", the second international joint conference on Autonomous agents and multi-agent systems Melbourne, Australia ACM Press, 2003.

About Authors

Gary Berg-Cross, Ph.D.

Dr. Berg-Cross is a Cognitive Psychologist (Ph.D. SUNY Stony Brook, BS RPI) with over 25 years of experience with ontological analysis growing out of his dissertation on knowledge representation for discourse understanding. He has taught human learning and cognition & problems solving courses at the University of Delaware, GW & GMU. Since 1975 he has consulted to health and medical facilities to evaluate and improve medical functioning and manage medical knowledge. Major thrusts of work includes reusable knowledge, vocabularies and semantic interoperability achieved through ontological analysis. For the last 7 years he has focused on the semantics of Enterprise Architecture and SOA models. Since 2005 he has been an active member of SICoP and facilitates workshop meetings including the June session on geospatial ontology that lead to the creation of a spatial community of practice (SOCoP). He is currently the SOCoP Executive Secretariat.

Augustine J. Kwon, Ph.D.

Dr. Kwon is a trained industrial engineer and has over 10 years of experience in modeling and simulation, artificial intelligence, and information system design, analysis, and development in various DoD projects. His research interest and areas of expertise include Simulation Modeling/Analysis, Agent based System Modeling, Applications of Neural Networks, information fusion and Semantic Reasoning in C2 environment, Probability models, and Operations Research. He is a senior associate at ICF International since August of 2006.

Wai-Tat Fu, Ph.D.

Dr. Fu is an applied cognitive scientist with interests in human–computer interaction, cognitive modeling, information seeking, interactive decision making, and cognitive skill acquisition; he is an Assistant Professor in the Human Factors Division and Beckman Institute of Advanced Science and Technology at the University of Illinois at Urbana–Champaign.