

Towards a Cognitive System for Decision Support in Cyber Operations

Alessandro Oltramari and Christian Lebiere
Functional Modeling Systems Lab
Department of Psychology
Carnegie Mellon University
Pittsburgh, USA

Wen Zhu
Alien Science and Technology
Washington D.C., USA

Lowell Vizenor
Refinery 29
New York, NY, USA

Randall Dipert
Department of Philosophy
University of Buffalo

Abstract— This paper presents the general requirements to build a “cognitive system for decision support”, capable of simulating defensive and offensive cyber operations. We aim to identify the key processes that mediate interactions between defenders, adversaries and the public, focusing on cognitive and ontological factors. We describe a controlled experimental phase where the system performance is assessed on a multi-purpose environment, which is a critical step towards enhancing situational awareness in cyber warfare.

Keywords—ontology, cognitive architecture, cyber security

I. INTRODUCTION

A cyber attack by a hostile nation-state or political organization is widely regarded as one of the most serious threats that the U.S. will face in the next decades. While greatly increased use of information systems has contributed enormously to economic growth, and has fueled a much more efficient and agile national defense, it has also made the U.S. enormously vulnerable to a variety of Internet and non-Internet cyber attacks, and to cyber espionage [1].

There are numerous factors that make cyber warfare and pure cyber defense, namely cyber security, especially problematic. The kinds of threats are diverse: destruction or theft of data, or interference with information systems and networks, across a spectrum of private and public interests. The legal and ethical status of cyber attacks or counterattacks by states are also unclear, at least when deaths or permanent destruction of physical objects does not result. It is still an open question what U.S. policy is or should be, and how cyber threats are analogous to traditional threats and policies—for example whether “first use” deterrence, and in-kind responses apply, and whether a policy of pure cyber defense does not put the far greater burden on attacked rather than attacking nations [2]. As this overview may suggest, untangling the complexity of cyber attacks becomes a key element for augmenting situational awareness in the cyber environment: in this position paper, we propose to tackle this problem from a semantic and cognitive modeling perspective, combining ontologies and

cognitive architectures into an intelligent system capable of supporting humans in cyber operations as well as acting autonomously as a team member.

The paper is divided into four main parts. After introducing some aspects of special interest to modeling cyber warfare (Section II), in Section III we present a hybrid decision support system based on cognitive architectures and ontologies. Section IV unfolds the experimentation plan to test the system by means of a scalable synthetic environment, and Section V delineates a framework of implementation centered on an object-based infrastructure.

II. RELEVANT CHARACTERISTICS OF CYBER WARFARE

In general, time variables play an important role in the design of decision support systems [3]: temporal constraints become even more stringent when those systems have to deal with cyber attacks, where real time responses are typically hindered by the knowledge-intensive nature of cyber operations and associated tasks. Some decisions on where and when to invoke various methods of cyber defense and mitigate damage, as well as decisions to launch a cyber counterattack, need to be made quickly. Large-scale cyber attacks or counterattacks are likely going to require careful, human decision-making for some time into the future. Yet there are other responses to cyber attacks or cyber espionage that could and should be done immediately, such as revoking an employee’s access if suspicious activity is detected, blocking all remote access or from certain URLs and through certain servers, immediate assessment of likely damage and risks, and so on. What we propose in this paper is the building of a cognitive system for decision support that will emulate ideal human responses to cyber attacks. This would be accomplished through careful design of its architecture, both in terms of cognitive mechanisms and knowledge resources, and by comparing its outputs on case studies with actions of human agents. The benefits are threefold. First, by cognitive modeling we come to better understand the mechanisms underlying human decisions in the realm of cyber warfare and cyber

This research was partially supported by a Defense Threat Reduction Agency (DTRA) grant number: HDTRA1-09-1-0053 to Christian Lebiere (Principal Investigator) and Alessandro Oltramari.

espionage, coupling the cognitive aspects and the semantic contents of decision-making. Second, after extensive testing we could use this intelligent decision-making system to recommend steps to human decision-makers—e.g., recommendations to gather further information, or actually to act in a certain way and to assess the risks of not acting. Finally, in cases where the reliability of the system is high, and where time is of the essence or the actions have little risk (such as revoking one employee’s system access, or access to one URL), the intelligent system could act swiftly and autonomously.

Some forms of attacks, such as Distributed Denial Of Service (DDoS) and other botnet jamming of networks or servers, show signs of admitting purely technological solutions. However, human error by employees has repeatedly been cited as the most common source of vulnerability [4], [5], [6], [7]. One technique of gaining illegitimate access to an information system that still appears with remarkable frequency is spear-phishing: emails to DOD employees or defense contractors with spoofed addresses from acquaintances that seem to have a harmless photograph, PDF, or other attachment¹. While this exploitation might not alone gain direct access to secure systems, it may allow an attacker to gather personal information that can be used to guess passwords, answer security questions, and so on. Social networking sites and other open data and the use of analytics allow attackers to identify employers, friends, relatives, shopping and driving habits, and so on. This aids an attacker enormously in the identification of targets and gaining access: for instance, in a recent case the New York Times’ sites were brought down when a group claiming to be the Syrian Free Electronic Army used social media and spear phishing to gain access to employees’ passwords to the server that handled the NY Times’ Domain Network System (DNS). Likewise even if smartphones and other portable devices are not used at secure locations and do not contain classified or sensitive data, hacking into them (or intercepting cellular and WiFi communications, including with vehicles and home monitoring devices) can provide personal data that can be utilized to make direct attacks.

III. TOWARDS A COGNITIVE SYSTEM FOR DECISION SUPPORT IN CYBER WARFARE

A. General methodology

Our approach is inspired by the notion of “sociotechnical system” [8], which emphasizes the interaction between people and technology in workplace. Ontology analysis has recently proved to be an effective tool for investigating these complex aspects [9]: nevertheless, the interactive nature of socio-technical systems demands a broader framework, where human behavior can be studied not only in terms of action schematics, planning and rules, but also as a genuinely cognitive phenomenon, which can be properly investigated only as a dynamic system. Accordingly, the key elements of our proposed method for modeling cyber operations are:

- Cognitive architecture – design and development of cognitive models of decision-making in cyber defense

¹ Because of their prevalence and complexity in terms of kind and number of cognitive agents, we intend to include these as paradigms of our use-cases.

based on ACT-R² cognitive architecture [10]. The models will focus on: learning mechanisms, memory and attentional limitations, decision-making strategies, risk perception, and trusted judgments.

- Ontologies – design and development of applied formal ontologies to 1) serve as a knowledge base for our cognitive models (*Cyber Security Ontologies*) and to 2) classify and annotate cyber security test and training data (*Scenario Ontologies*).
- Live, Virtual, Constructive (LVC) Integration – Enable the analysis of cyber defense strategies; support training for cyber security personnel; validate the cognitive models developed with an attack/mitigate/counter-attack scenarios and enhance them by leveraging learning mechanisms.

By integrating these elements in a coherent multi-purpose system, we aim at unraveling the complex structures that mediate interactions among defenders, adversaries and the public: in this respect, the overall goal is to enhance situational awareness in cyber warfare by assessing human performance in a simulated environment. The system is also meant to interact autonomously in a *hybrid* team, i.e. playing the role of a “teammate” sentinel in support of humans, eventually capable of prompting decisions and perform actions in more mature stages of development.

To provide a richer characterization of our approach, Section B illustrates the functional requirements of the envisioned system, while Section C and D will narrow the focus to, respectively, ACT-R cognitive architecture (the central component of the system) and the ontologies needed to frame the knowledge component of the architecture.

B. Functional models of cyber operations

Modeling decision-making in the cyber security framework requires multiple factors to be investigated: (i) the size and the variety of *knowledge* which is necessary to classify and analyze attacks and defensive actions; (ii) the *flexible behavior* required by coupling alternative strategies of response to specific cyber threats, updating and revising strategies when the circumstances of the attack or the environmental conditions evolve; (iii) *learning by experience* how to deal with cyber attacks; (iv) *interacting* in a team by building a mental representation of the co-workers as well as of the enemies. These factors can be mapped to the 12 criteria distilled in [11] (from the original list compiled by Newell in [12]) that a cognitive architecture would have to satisfy in order to achieve human-level functionality. In these regards, cognition is not considered as a “tool” for optimal problem solving but, rather, as a set of limited information processing capacities (so-called ‘bounded rationality’ [13])³. In a similar fashion, Wooldridge had identified the requirements that an agent should satisfy in order to act on a rational basis [14], namely: *reactivity*, the capacity of properly reacting to perceptual stimuli; *proactivity*, the capacity of operating to pursue a goal; *autonomy*, implying

² Pronounced, “act-ARE”: Adaptive Control of Thought—Rational.

³ Despite the relevance of emotions in decision-making [34], our approach doesn’t extend to the investigation of affective aspects at this stage.

an unsupervised decision making process; *social ability*, the capacity of interacting with other agents and revising mental states accordingly.

State-of-the-art research on cognitive architectures (SOAR, ACT-R, CLARION, OpenCog, LIDA, etc.) has produced a significant amount of results on specifying this extensive range of functions⁴: by and large, ACT-R has accounted for the broadest range of cognitive activities at a high level of fidelity, reproducing aspects of human data such as learning, errors, latencies, eye movements and patterns of brain activity [10]. However, these results have often involved relatively narrow and predictable tasks. Most importantly, cognitive architectures have just started to tackle the problem of how to model *social ability* [15], which is a crucial aspect of our approach. A fundamental feature of human social ability is “mindreading” [16], i.e. to understand and predict the actions of others by means of postulating their intentions, goals and expectations: this process of interpretation is feasible only if an agent can learn to *represent* the mental states of others on the basis of cumulative experience and background knowledge, combining the resulting mental model with the continuous stream of data from the environment, aiming at replicating the cognitive processes that have likely motivated the other agents to perform the observed actions. Scaling up ACT-R to account for more extensive multi-agent scenarios can help to build comprehensive models⁵ of social conflict and cooperation, which are critical to discern the governing dynamics of cyber defense. But if leveraging the ACT-R framework might be sufficient to replicate the *mechanisms* described in (ii)-(iv), the knowledge functionality (i) can to be fulfilled only by injecting a fair amount of highly expressive knowledge structures into the architecture: accordingly, ontologies can provide these structures in the form of semantic specifications of declarative memory *contents* [17]. As [18], [19], and [20] show, up to this time most research efforts have focused on designing methods for mapping large knowledge bases to ACT-R declarative module, but with scarce success. Here we commit to a more efficient approach: modular ontologies. Modularity has become a key issue in ontology engineering. Research into aspects of ontological modularity covers a wide spectrum: [21] gives a good overview of the breadth of this field. Our modular approach guarantees wide coverage and “manageability”: instead of tying ACT-R to a single large ontology, which is hard to maintain, update and query, we propose a *suite* of ontologies that reliably combine different dimensions of the cyber defense context, e.g. representation of secure information systems at different levels of granularity (requirements, guidelines, functions, implementation steps); categorization of attacks, viruses, malware, worms, bots; descriptions of defense strategies; the mental attitudes of the assailant, and so on.

In our context, the computational system resulting from the combination of cognitive and knowledge functionalities aims at fostering a better understanding of cyber attacks, supporting human operators in cyber warfare, eventually cooperating with

them in well-defined synthetic environments. The rest of the paper presents in more detail the basic components of such a hybrid framework.

C. Replicating cognitive mechanisms with ACT-R

Cognitive architectures attempt to capture at the computational level the invariant mechanisms of human cognition, including those underlying the functions of control, learning, memory, adaptivity, perception, decision-making, and action. ACT-R [10] is a modular architecture including perceptual, motor and declarative memory components, synchronized by a procedural module through limited capacity buffers (see figure 1 for the general diagram of the architecture). Declarative memory module (DM) plays an important role in the ACT-R system. At the symbolic level, ACT-R agents perform two major operations on DM: 1) accumulating knowledge “chunks” learned from internal operations or from interacting with objects and other agents populating the environment and 2) retrieving chunks that provide needed information. ACT-R distinguishes ‘declarative knowledge’ from ‘procedural knowledge’, the latter being conceived as a set of procedures (production rules or “productions”) which coordinate information processing between its various modules [10]: according to this framework, agents accomplish their goals on the basis of declarative representations elaborated through procedural steps (in the form of *if-then* clauses). This dissociation between declarative and procedural knowledge is grounded in experimental cognitive psychology; major studies in cognitive neuroscience also indicate a specific role of the hippocampus in “forming permanent declarative memories” and of the basal ganglia in production processes (see [22], pp. 96-99, for a general mapping of ACT-R modules and buffers to brain areas and [23] for a detailed neural model of the basal ganglia’s role in controlling information flow between cortical regions). ACT-R performs cognitive tasks by combining rules and knowledge: for reasons of space, a complete analysis of how the architecture instantiates this cognitive-based processing is not suitable here. Nevertheless, two core mechanisms need to be mentioned: *i) partial matching*, the probability of association between two distinct declarative knowledge chunks, computed on the basis of adequate similarity measures (e.g. a bag is more likely to resemble a basket than a tree); *ii) spreading of activation*, the phenomenon by which a chunk distributionally activates the different contexts in which it occurs (a bag can evoke shopping, travel, work, etc.). These two basic mechanisms belong to the general sub-symbolic computation underlying chunk activation, which in ACT-R controls the retrieval of declarative knowledge elements by procedural rules. In particular, ACT-R chunk activation is calculated by the following equation:

$$A_i = \ln \sum_j t_j^{-d} + \sum_k W_k S_{ki} + \sum_l MP_l Sim_{li} + N(0, \sigma) \quad (1)$$

On the basis of the first term, the more recently and frequently a chunk *i* has been retrieved, the higher the activation and the chances of being retrieved (t_j is the time elapsed since the j^{th} reference to chunk *i* and d represents the memory decay rate). In the second term of the equation, the contextual activation of a chunk *i* is set by the attentional

⁴ See [33] for a comprehensive overview of the most recent advancements in the area of cognitive architectures research.

⁵ Note that the distinction between ‘model’ and ‘agent’ when dealing with cognitive architectures is a blurred one. For clarity’s sake we will henceforth use ‘agent’ to avoid ambiguities with the notion of semantic model (ontology). In general, an agent is a cognitive model that dynamically interacts with the environment.

weight W_k , given the element k and the strength of association S_{ki} between k and the i . The third term states that, under partial matching, ACT-R can retrieve the chunk that matches the retrieval constraints to the greatest degree, combining the similarity Sim_{li} between l and i (a negative score that is assigned to discriminate the ‘distance’ between two terms) with the scaling mismatch penalty MP . The final factor of the equation adds a random component to the retrieval process by including Gaussian noise to make retrieval probabilistic.

The intertwined connection between declarative and procedural knowledge, weighted by stochastic computations, represents the necessary substrate for realizing at the computational level the functionalities outlined in section B: more specifically, we claim that ACT-R can successfully be used to emulate human behavior in selecting and executing defense strategies, matching input data from on-going cyber attacks to deeply structured background knowledge of cyber operations. In the past, ACT-R architecture has been successfully used in context where integrating declarative and procedural knowledge was also a fundamental issue, e.g. air traffic control simulations [24].

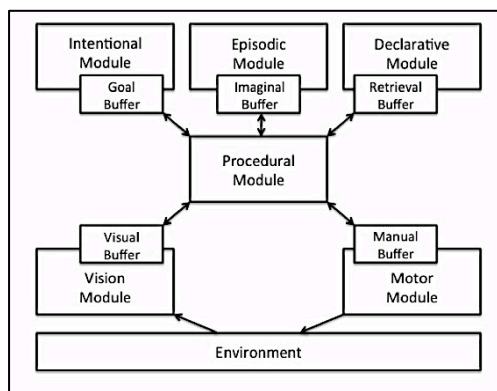


Figure 1 ACT-R Modular Structures

D. Augmenting ACT-R with cyber security ontologies

The development of cyber security ontologies is a critical step in the transformation of cyber security from an art to a science. In 2010, the DOD sponsored a study to examine the theory and practice of cyber security, and evaluate whether there are underlying fundamental principles that would make it possible to adopt a more scientific approach. The study team concluded that:

The most important attributes would be the construction of a common language and a set of basic concepts about which the security community can develop a shared understanding. A common language and agreed-upon experimental protocols will facilitate the testing of hypotheses and validation of concepts [25].

The need for controlled vocabularies, taxonomies, and ontologies to make progress toward a science of cyber security is recognized in [26] and [27] as well. In the domain of cyber security, the ontologies would include, among other things, the classification of cyber attacks, cyber incidents, and malicious and impacted software programs. From our point of

view, which seeks to accurately represent the human-side of cyber security, we also expand our analysis to: (i) the different roles that system users, defenders and policy makers play in the context of cyber security; (ii) the different jobs and functions that the members of cyber defender team play and the knowledge, skills and abilities needed to fulfill these functions. In order to reduce the level of effort, we will reuse existing ontologies when possible⁶ and only create new ontologies that support the use cases we select.

The decentralization of knowledge organization and maintenance to a variety of interconnected ontology modules leverages a shared bridging component, i.e. BFO reference ontology⁷: in this sense, BFO plays the role of the common semantic infrastructure to define, populate and update multiple context-driven cyber ontologies. The various modules will be encoded in W3C language OWL⁸: the process of porting them into ACT-R is managed automatically at the architecture level by built-in LISP functions, which are able to *a.* read and interpret the XML-based syntax of the semantic model and *b.* convert it into ACT-R declarative format. A set of broad schemas drives this conversion process: for instance, the direct mapping between the ‘‘chunk-type’’ primitives in ACT-R and classes in the ontologies has been designed. Further schemas at a narrower level of granularity will be provided, as engineered for an analogous framework presented at STIDS 2012 [28].

IV. COGNITIVE SIMULATIONS OF CYBER OPERATIONS

A. Experimental Design

The first objective of building an intelligent system endowed with adequate representation of cyber security knowledge is to use it in scalable synthetic environments for training human decision-makers. In addition, once the system has incorporated the necessary rational capabilities (defined in the previous section) and learned the dynamics of team interaction, we aim at testing the possibility of deploying it as an autonomous defensive agent in virtual cyber operations. In order to achieve the necessary degree of robustness and dependability, we plan simulations at different levels of complexity, as follows:

BSE — *Basic Synthetic Environment*: two ACT-R agents face each other playing the role of assailant and defender;

HSE — *Hybrid Synthetic Environment*: an ACT-R agent and a human face each other playing the role of assailant and defender;

HSGE — *Hybrid Synthetic Group Environment*: two teams, each constituted by humans and ACT-R agents face each other playing the role of assailant and defender.

In order to run these incremental simulations, we will initially collect an experimental dataset of cyber attacks, to be split into train and test set. In particular, we will focus on spear phishing attacks, as delineated in section II. The datasets will be organized to instantiate classes and properties of the defined modular ontologies. Each level of the cognitive-based simulation will be conceived as a block composed of multiple

⁶ For instance, exploiting material from this portal: <http://militaryontology.com/cyber-security-ontology.html>

⁷ <http://ontology.buffalo.edu/bfo/>

⁸ <http://www.w3.org/TR/owl-features/>

trials⁹. At the **BSE** level, the simulation aims at assessing the soundness of the cognitive mechanisms executed by the agent, serving also as a system debugging and evaluation of experimental settings. In the **HSE**, the agent will have to compete against humans, whose potentially erratic behavior will be exploited by the agent as a primary source of acquisition of cyber warfare strategies and mental representation of the opponent. Finally, in **HSGE** the scenario will get more complex by shifting to a multi-agent framework, where each defending agent will have to learn intra-group cooperation and build mental representation of the opponent as a group (whose members act complementarily and collectively to harm the defending team).

In the delineated experimental phase we plan to expand our previous work on applying cognitive architectures to decision-making in non-zero sum games [29]: cooperative and conflicting phenomena have been comprehensively studied using game theory [30], in which complex social dynamics are narrowed down to relatively simplified frameworks of strategic interaction. Valid models of real-world phenomena can provide better understanding of the underlying socio-cognitive variables that influence strategic interaction: of course these models need be consistent with the structural characteristics of games, and with the actual everyday situations at hand. In this respect, the goal of the planned cognitive simulations is to study decision-making by deploying computational rational agents in cyber attack “gamified” scenarios.

B. Evaluation plans

As recent studies have shown [31], training users to respond to cyber attacks becomes effective only after several iterations. But high time-costs in training can expose socio-technical systems to harmful consequences, with no chance of recovering stolen information or, even worse, of fully restoring the functionalities of the system. Our approach aims to improve cyber defense strategies and speed up the deployment of counter-measures. In particular, we plan to assess the correspondence between the models’ simulations and the human behavior in cyber-operations by analyzing human data in decision-making processes. Accordingly, we will apply different analytical methods, such as computing means and standard errors (for decisions), medians and the 1st and 3rd quartiles (for decision times) — similar approaches have been successfully proposed in [32]. We will encode conversion functions in the system to format the outputs as discrete decisions (e.g. “delete spear phishing email”, “scan for malware”, “reactivate firewall”, etc.). Exploiting ACT-R internal clock module, we will also be able to reproduce decision times at human granularity scale, tracking the relevant stages of the rational decision-making process.

V. APPLICATION FRAMEWORK

So far we have discussed the general requirements and described the high-level cognitive structures of an intelligent system for decision support in cyber warfare. However, a product or a solution based on these requirements and architecture will need to address specific problems in the business domain. Furthermore, the end product would likely

require integration with other technical components and frameworks. We see an opportunity to apply the concepts described in this paper for the development of an application capable of assessing and reducing information systems vulnerabilities through live, virtual, and constructive (LVC) simulations. Such an application can support a wide range of cyber defense objectives, including: (i) analysis of cyber defense strategies and identification of network vulnerabilities through simulated attacker-defender interaction in BSE – HSE – HSGE scenarios; (ii) training for cyber security personnel with a suitable ACT-R agent simulating the attacker against human players; (iii) validation and enhancement of the cognitive models developed with an attack counter-attack scenario. To support LVC simulations, the application will need to work with existing distributed modeling and simulation infrastructures, such as the High Level Architecture (HLA)¹⁰ or Testing and Training Enabling Architecture (TENA)¹¹. The key integration activities include:

- Identification and creation of reusable ‘objects’. A distributed modeling and simulation framework such as TENA encourages objects representing things such as targets and assets to be reused across simulations¹². In particular, within the intelligent decision support system, we see opportunities at two levels: 1) creation of reusable objects representing attackers and defenders (these objects can be used to simulate behaviors of the actors); 2) creation of reusable objects representing IT Infrastructure components that could be under cyber attacks (these objects model the commands and instructions that can be sent to various components and their responses).
- Integration of reusable objects in to the middleware layer of the modeling and simulation framework. Figure 2 shows a reusable TENA object (representing cyber attackers) plugged into the middleware layer.
- Implementation of runtime knowledge sharing in the modeling and simulation framework. In the example shown in figure 2, the ACT-R cognitive model (representing the defender) is integrated with knowledge sources incrementally stored in ACT-R declarative memory module: a) modular cyber security ontologies, retrieved from the TENA Repository and; b) the modular ontologies of the **scenario [1]**, incrementally stored in TENA Event Data Management.

¹⁰ <http://standards.ieee.org/findstds/standard/1516-2010.html>

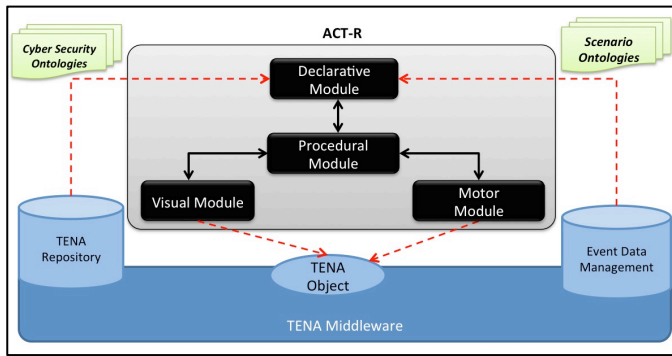
¹¹ Test and Training Enabling Architecture (TENA): <https://www.tena-sda.org/display/intro/Home>

¹² TENA object-oriented modeling features well fit our ontology-driven cognitive system.

⁹ Setting to 100 the number of trials should guarantee a satisfactory level of stochasticity in the results.

VI. CONCLUSION

The novelty of our approach relies on grounding a decision support system in a broad spectrum of human-level cognitive functionalities blended with highly structured knowledge resources. In particular, by focusing on learning mechanism, context-driven semantic specifications and scalable simulations, the obtained computational system can serve both as a training environment for cyber personnel and as autonomous team member operating in advanced security settings. Our position paper aims at fostering the discussion within the communities of interest and can play the role of a starting platform for a scientific project proposal.



[2]

Figure 2 The Cognitive System realized in the TENA framework.

REFERENCES

- [1] R. Dipert, "Other-Than-Internet(OTI) Cyberwarfare: Challenges For Ethics, Law, and Policy," *Journal of Military Ethics*, vol. 12, no. 1, pp. 34-53, April 2013.
- [2] R. R. Dipert, "The Ethics of Cyberwarfare," *Journal of Military Ethics*, vol. 9, no. 4, pp. 384-410, December 2010.
- [3] M.I. Hwang, "Decision Making under time pressure: A model for information systems research," *Information and Management*, vol. 27, pp. 197-203, 1994.
- [4] Symantec. (2013, June) Symantec. [Online]. HYPERLINK "http://www.symantec.com/about/news/release/article.jsp?prid=20130605_01"
http://www.symantec.com/about/news/release/article.jsp?prid=20130605_01
- [5] Brian Montopoli. (2013, May) CBS News. [Online]. HYPERLINK "http://www.cbsnews.com/8301-201_162-57586624/how-chinese-hackers-steal-u.s-secrets/"
http://www.cbsnews.com/8301-201_162-57586624/how-chinese-hackers-steal-u.s-secrets/
- [6] (2013, May) Wall Street Journal. [Online]. HYPERLINK "http://online.wsj.com/article/PR-CO-20130530-906764.html?mod=googlenews_wsj"
http://online.wsj.com/article/PR-CO-20130530-906764.html?mod=googlenews_wsj
- [7] J. C. Forsythe, A. Silva, S. Stevens-Adams, and J. Bradshaw, "Human Dimensions in Cyber Operations Research and Development Priorities," SANDIA Report 2012-9188, Technical 2012.
- [8] K. B. De Greene, *Sociotechnical systems: factors in analysis, design, and management.*: Prentice-Hall, 1973.
- [9] N. Guarino, E. Bottazzi, R. Ferrario, and G. Sartor, "Open Ontology-Driven Sociotechnical Systems: Transparency as a Key for Business Resiliency," in *Information Systems: Crossroads for Organization, Management, Accounting and Engineering.*, 2012, pp. 535-542.
- [10] John R. Anderson and Christian J Lebiere, *The Atomic Components of Thought.*: Erlbaum, 1998.
- [11] John R. Anderson and Christian Lebiere, "The Newell Test for a theory of cognition," *Behavioral and Brain Sciences*, vol. 26, no. 5, pp. 587-637, 2003.
- [12] Allen Newell, *Unified Theories of Cognition*. Cambridge, Massachusetts: Harvard University Press, 1990.
- [13] H. Simon, "Bounded Rationality and Organizational Learning," *Organization Science*, vol. 2, no. 1, pp. 125-134.
- [14] M. Wooldridge, *Reasoning about Rational Agents*. Cambridge, MA, United States of America: The MIT Press, 2000.
- [15] R. Sun, *Cognition and Multi-agent Interaction*, R. Sun, Ed.: Cambridge University Press, 2006.
- [16] Paul Bello, "Cognitive Foundations for a Computational Theory of Mindreading," *Advances in Cognitive Systems*, vol. 1, pp. 59-72, 2012.
- [17] A. Oltramari and C. Lebiere, "Knowledge in Action: Integrating Cognitive Architectures and Ontologies," in *New Trends of Research in Ontologies and Lexical Resources*, Alessandro, Vossen, Piek Oltramari, Lu Qin, and Ed. Hovy, Eds.: Springer, pp. 135-154.
- [18] J. Ball, S. Rodgers, and K. Gluck, "Integrating ACT-R and Cyc in a large-scale model of language comprehension for use in intelligent agents," in *Papers from the AAIL Workshop*, Menlo Park, CA, pp. 19-25.
- [19] B. J. Best, N. Gerhart, and C. Lebiere, "Extracting the Ontological Structure of OpenCyc for Reuse and Portability of Cognitive Models.," in *Proceedings of the 17th Conference on Behavioral Representation in Modeling and Simulation*, 2010.
- [20] S. Douglas, J. Ball, and S. Rodgers, "Large declarative memories in ACT-R," in *Proceedings of the 9th International Conference of Cognitive Modeling*, Manchester, UK.
- [21] H. Stuckenschmidt, C. Parent, and S. Spaccapietra, "Modular Ontologies - Concepts, Theories and Techniques for Knowledge Modularization," , 2009.
- [22] J. R. Anderson, *How Can the Human Mind Occur in the Physical Universe?* New York: Oxford University Press.
- [23] A. Stocco, C. Lebiere, and J. R. Anderson, "Conditional Routing of Information to the Cortex: A Model of the Basal Ganglia's Role in Cognitive Coordination," *Psychological Review*, vol. 117, no. 2, pp.

541-574, 2010.

- [24] C. Lebiere, "Constrained Functionality: Application of the ACT-R Cognitive Architecture to the AMBR Modeling Comparison." Mahwah, NJ: Erlbaum, 2005.
- [25] The MITRE Corporation, "Science of Cyber-Security," The MITRE Corporation, McLean, VA, Technical 2010 (extract).
- [26] D. A. Mundie and D. M. McIntire, "The MAL: A Malware Analysis Lexicon," CERT® Program - Carnegie Mellon University, Technical 2013.
- [27] Randall Dipert, "The Essential Features of an Ontology for Cyberwarfare," in *Conflict and Cooperation in Cyberspace - The Challenge to National Security*, Panayotis A Yannakogeorgos and A. B. Lowther, Eds.: Taylor & Francis, 2013, pp. 35-48.
- [28] A. Oltramari and C. Lebiere, "Using Ontologies in a Cognitive-Grounded System: Automatic Action Recognition in Video Surveillance," in *Proceedings of STIDS 2012 (7th International Conference on "Semantic Technology for Intelligence, Defense, and Security")*, Fairfax, VA, 2012.
- [29] A. Oltramari, C. Lebiere, N. Ben-Asher, and C. Gonzalez, "Strategic Dynamics Under Alternative Information Conditions," in *Proceedings of ICCM 2013 (International Conference of Cognitive Modeling)*, Ottawa, 2013.
- [30] A. Rapoport, M. J. Guyer, and D. G. Gordon, *The 2 x 2 game*. Ann Arbor, MI: University of Michigan Press, 1976.
- [31] B. M. Bowen, D. Ramaswamy, and S. Stolfo, "Measuring the Human Factor of Cyber Security," *Homeland Security Affairs*, vol. 5, no. 2, 2012.
- [32] J. N. Marewski and K. Mehlhorn, "Using the ACT-R Architecture to specify 39 quantitative process models of decision making," *Judgement and Decision Making*, vol. 6, pp. 439-519, August 2011
- [33] J.E. Laird, *The SOAR Cognitive Architecture*. USA: The MIT Press, 2012.
- [34] C. L. Dancy, F. E. Ritter, and F. E. Berry, "Towards adding a physiological substrate to ACT-R," in *Proceedings of the 21st Conference on Behavior Representation in Modeling and Simulation*, Amelia Island, FL, 2012, pp. 78-85.