Human, Model and Machine: A Complementary Approach to Big Data

Robert Thomson
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA, 15213
1-412-268-6284
thomsonr@andrew.cmu.edu

Christian Lebiere
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA, 15213
1-412-268-6028
cl@cmu.edu

Stefano Bennati
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA, 15213
1-412-268-6284
sbennati@andrew.cmu.edu

ABSTRACT
In this paper, we describe a framework for processing big data that maximizing the efficiency of human data scientists by having them primarily operate over information that is best structured to human processing demands. We accomplish this through the use of cognitive models as an intermediary between machine learning algorithms and human data scientists. The ACT-R cognitive architecture is a computational implementation of a unified theory of cognition. ACT-R cognitive models can take weakly structured data and learn to filter information and make accurate inferences orders of magnitude faster than machine learning, and then present these well-structured inferences to human data scientists. The role for human data scientists is both oversight and feedback; one complementary piece of a hierarchy of cognitive and machine learning techniques that are computationally appropriate for their level of information complexity.

Categories and Subject Descriptors
A.0 [**General**]: Conference Proceedings

General Terms
Algorithms, Theory.

Keywords
Cognitive Architectures, ACT-R, Deep Learning, Big Data

1. BACKGROUND
The effective and efficient analysis of big data requires the complementary engagement of human data scientists, cognitive models, and machine learning algorithms. Each component brings together distinct yet complementary characteristics. Similar to the DIKW (*Data-Information-Knowledge-Wisdom*) [1] cycle of big data research, we argue for a pyramidal structure of data refinement, analogous to the process of sensemaking, albeit with more bidirectional interconnections than the bottom-up DIKW approach (see Figure 1).

HCBDR '14, April 01 - 03 2014, Raleigh, NC, USA

ACM 978-1-4503-2938-5/14/04…$15.00.
http://dx.doi.org/10.1145/2609876.2609883

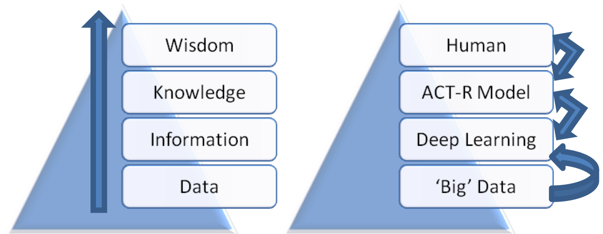

Figure 1. Sample knowledge tree where information becomes more structured as it reaches progressively higher stages of the pyramid.

The DIKW pyramid (left) begins with large amounts of unstructured (*big*) *data* at its base. By filtering and extracting relevant features from the raw data, a primitive semantics is developed, which transitions data to *information* [16]. By contextualizing and hierarchically organizing this information, we generate *knowledge*. This knowledge is akin to the kind of sensemaking frames seen in intelligence analysis [8]. Finally, *wisdom* is gained when knowledge is transformed into actionable elements with refined semantics (Koomey, 2008). By analogy, we propose a data analysis hierarchy with massively-parallel *machine learning* algorithms operating over unstructured data at its base to filter and organize them, *cognitive models* refining this loosely structured information with more elaborate knowledge semantics, and finally *human analysts* at the top level providing overall awareness of context in organizing knowledge into actionable intelligence. Each processing layer matches the characteristics of the data layer that it transforms [21].

2. DATA PYRAMIDS
This section will present our analogy to the traditional DIKW pyramid in detail, describing the nature of each layer and their integration and interactions for both processing and learning.

2.1 Information / Machine Learning
Machine Learning algorithms like Deep Learning are massively parallel and can cope with large amounts of data; however they are limited because of relatively primitive semantics and thus are best suited for the initial filtering and structuring of data. Machine learning algorithms provide a scalable complement to the human analyst by structuring large amounts unfiltered data into information. To assume that machine learning can replace data scientists entirely is to ignore their current inability to capture the complex contextual semantics needed for high-level inference in domains such as computer vision and the daunting combinatorics of achieving data coverage in real-world intelligence scenarios.

Computer vision, by analogy, is typically viewed as a bottom-up process refining raw inputs like pixels into increasingly high-level features. However, the process has proven so difficult it is considered unsolvable due to the ambiguities and combinatorics in raw input. The emerging solution is to view vision as a bi-directional process that utilizes top-down input in the form of contextual expectations to provide strong priors for constraining bottom-up perception. This view is compatible with the evolving understanding of the neuroscience of human vision, where what was initially thought of as a bottom up process of generating increasingly abstract representations is how understood to be deeply influenced by top-down expectations for higher-level cognitive processes. The same criticisms have been applied to sensemaking, where attempts to cull and organize data in a bottom-up, context-free way has led to countless intelligence failures [14]. Our solution to this problem is to use cognitive models, constrained by an architecture (in our case, ACT-R), to provide not only these top-down expectations for bottom-up perceptual processes, but to also operate independently over these data and structure them in a way that is more appropriate for the capabilities and limitations of human data scientists.

2.2 Knowledge / Cognitive Models

Cognitive models are capable of transforming the loosely-structured data coming from machine-learning algorithms into a more structured form appropriate for human data scientists. We argue that the best way for models to do this is to be grounded in a cognitive architecture. A cognitive architecture is a set of computational processes representing cognitive processes including attention, memory and pattern-matching, organized in a structure reflecting the organization of the human brain. A cognitive architecture provides a principled platform for developing behavior models constrained by human processes and predictive of human performance, but also with the means to mitigate limitations in human reasoning such as cognitive biases. Cognitive models – built within the constraints of the architecture – provide a computational implementation of cognitive mechanisms that can be de-biased through the adoption of proper procedures such as Structured Analytic Techniques and scaled up beyond human abilities such as through faster processing time and larger memory. Still, by matching human expert patterns against new situations, they transform information into knowledge.

An advantage of using a cognitive model as a smart filter is that the underling architecture provides some of the benefits of human heuristic reasoning while allowing for more control over biases in reasoning. The model relieves human data scientists of the more time-consuming pattern-matching effort required in data wrangling, while also providing a more bias-free implementation of that functionality. In the IARPA ICArUS (Integrated Cognitive-Neuroscience Architectures for Understanding Sensemaking) project, we have developed an integrated neurocognitive model of complex sensemaking processes that was capable of matching human performance and predicting (and mitigating) the sources of multiple biases [8]. Cognitive biases arise from the interaction between (a) powerful but limited cognitive processes, (b) complex, demanding sensemaking tasks in a dynamic fast-paced information-rich environment, and (c) adaptive heuristic strategies adopted to accomplish the tasks under strong time constraints using limited means. The notion of architecture is essential to this argument. It is because of the constraints of the cognitive architecture that we can understand the sources of biases in analysts, and provide alternative approaches in the form of cognitive models that alleviate them.

Using the notion of architecture, it is also appropriate to view our processing hierarchy as an integrated architecture with complex ***data flows*** including both top-down expectations and bottom-up filtering and elaboration, and corresponding ***learning flows*** that continually tune the system for improved performance. This stands in contrast to the predominantly bottom-up processing of the DIKW pyramid, which we argue falls prey to the same limitations of solely bottom-up processing approaches as in other fields. In another analogy from the memory literature, the process of learning is called a process of *consolidation* and *elaboration*, representing the fact that organizing data in memory is an organic process rather than a set of sequential filters.

This is not to argue that cognitive architectures are a panacea for data analysis research, or are 'only one breakthrough away' from replacing data analysts. While recent advances in cognitive architectures are exploring novel brain-inspired mechanisms to create hierarchies of information [20], we still require human feedback to build increasingly abstract and semantically rich levels of structured representations. Hierarchical knowledge structures – by compartmentalizing the data space – make decision-making more efficient and are an essential component of applying cognitive methods to big data.

2.3 Wisdom / Human Data Scientists

Cognitive models operate over a limited range of semantics and require human specification, whether through feedback/oversight or in the model design process itself. Data scientists have extremely rich semantics but very limited capacity, thus their place is at the top of the pyramid providing the final step: they provide cognitive oversight over the whole system and as such, they transform knowledge into wisdom. This process of oversight provides feedback to the model, providing additional learning instances to the cognitive model from which to make future decisions, and feedback to deep learning over the usefulness of learned features. Learning flows reciprocal to data flows are an essential component of an adaptive intelligent architecture.

Data scientists can elaborate this richly structured knowledge into a set of actionable elements by leveraging the strengths of human cognition operating over a much more constrained decision space. By setting humans in the complementary role at the tip of the hierarchy, we are leveraging their conceptual inference strengths and limiting the man-hour requirements that go into data wrangling [2]. The notion of a human as an overseer and capstone over a combination of machine learning and cognitive models has been applied in other domains. In the Mind's Eye project [15] concept of a smart camera, machine learning algorithms extracted perceptual features from a camera input, and a cognitive model operated over those visual features to match them to known patterns and detect suspicious behaviors, which are then flagged for a human security officer to examine. This feedback is then filtered back into the cognitive model in the form of new knowledge patterns and into the perceptual algorithms as feature utility.

Of great concern to data scientists (similar to intelligence analysts) is the influence of cognitive biases. The massive amounts of information that data scientists must sift through make an ideal environment for cognitive biases such as confirmation bias and the availability heuristic. In each case, the inability of human working memory to maintain a sufficient trace of all incoming information leads to heuristic shortcuts being used and improper generalizations being made. In the case of confirmation bias an analyst may only search for information that supports their current hypothesis while not looking for evidence that would disconfirm it. In the case of the availability heuristic underlying a satisfaction-of-search bias, the analyst may cease to look for alternative explanations once an initial hypothesis is formed. These are natural, and often effective, ways of coping with increasing demands on a fixed cognitive capacity. What is needed is an external augmentation of those capabilities that integrate naturally because they work on similar principles but within their limitations.

3. ACT-R COGNITIVE ARCHITECTURE

Before presenting examples of the efficacy of cognitive models, we will describe the ACT-R cognitive architecture. ACT-R is a computational implementation of a unified theory of cognition. It accounts for information processing in the mind via task-invariant mechanisms constrained by the biological limitations of the brain [2]. While sensemaking theory abstracts away from brain processes, it makes commitments to the control and flow of information that are commensurable with ACT-R's functional perspective. For example, the elaboration and reframing loops in sensemaking can be instantiated in the production rules controlling the flow of control and information in ACT-R. Furthermore, ACT-R is committed to localization of neural architecture, allowing for functional models to guide the development of neurally-inspired models. As such, a cognitive architecture like ACT-R provides a conceptual bridge between coarse-grained naturalistic processes such as sensemaking and fine-grained neural-level computational constructs.

The ACT-R architecture is organized as a set of modules, each devoted to processing a particular kind of information, which are integrated and coordinated through a centralized production system module (see Figure 2). Each module is assumed to access and deposit information into buffers associated with the module, and the central production system can only respond to the contents of the buffers, not the internal encapsulated processing of the modules. For instance, the goal module stores and retrieves information that represents the internal intention and problem solving state of the system and provides local coherence to behavior. These information bottlenecks correspond to important attentional and working memory limitations of the architecture but they enable scalable (i.e., parallelizable), tractable processes in the corresponding modules.

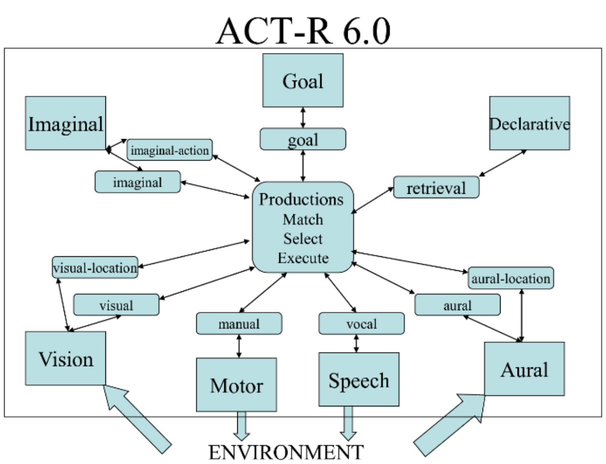

Figure 2. An overview of ACT-R's modules and their dependent buffers.

The declarative memory and production system modules, respectively, store and retrieve information that corresponds to declarative knowledge and procedural knowledge. Declarative knowledge is the kind of knowledge that a person can attend to, reflect upon, and usually articulate in some way, while procedural knowledge consists of the skills we display in our behavior, generally without conscious awareness. This distinction between explicit and implicit knowledge provides a powerful model of the human ability to acquire information from its environment and continually improve its performance.

Declarative knowledge in ACT-R is represented formally in terms of chunks, which corresponds to the episodic and semantic knowledge that promotes long-term coherence in behavior. Chunks have an explicit type, and consist of a set of slot-value pairs of information. Chunks are retrieved from declarative memory (DM) by an activation process. When a retrieval request is made the most active matching chunk is returned, where activation is computed as the sum of base-level activation, spreading activation, mismatch penalty and stochastic noise. Base-level activation reflects a chunk's recency and frequency of occurrence. Activation spreads from the current focus of attention through associations among chunks in declarative memory. These associations are built up from experience, and reflect how chunks co-occur in cognitive processing. Chunks are also compared to the desired retrieval pattern using a partial matching mechanism that subtracts from the activation of a chunk its degree of mismatch, additive for each component of the pattern and corresponding chunk value. Finally, noise is added to chunk activations to make retrieval a probabilistic process governed by a Boltzmann (softmax) distribution. These subsymbolic processes endow the cognitive architecture with much of the graded and adaptive nature of human cognition.

While the most active chunk is usually retrieved, a blending process can also be applied that returns a derived output reflecting the similarity between the values of the contents of all chunks, weighted by their retrieval probabilities reflecting their activations [9]. This blending process is used intensively in models of decision-making since it provides a tractable way to generalize decisions in continuous domains such as probability space.

The flow of information is controlled in ACT-R by a production system, which operates on the contents of the buffers. Each production consists of if-then condition-action pairs. Conditions are typically criteria for buffer matches, while the actions are typically changes to the contents of buffers that might trigger operations in the associated modules. The production with the highest utility is selected to fire from among the eligible productions. This control level enables the learning and execution of complex information processing strategies by the architecture itself to model the human ability to learn to perform new tasks almost without limitation.

3.1 Instance-Based Learning

Instance-based learning theory (IBL) is the claim that implicit expertise is gained through the accumulation and recognition of experienced events or instances. IBL was formulated within the principles and mechanisms of cognition in ACT-R, and makes use of the dynamics of chunk retrieval and blended retrievals [6]. The main claim of IBL is that implicit knowledge is generated through the creation of instances. These instances are represented in chunks with slots containing the conditions (e.g., a set of contextual cues), the decision made (e.g., an action), and the outcome of the decision (e.g., the utility of the decision).

IBL offers constraints on explanation by grounding implicit learning within the mechanisms of a cognitive architecture. For instance, the dynamics of an instance's sub-symbolic activations (e.g., frequency and recency in the base-level activation equation) provide a scientifically-justified mechanism for determining which instances are likely to be retrieved for a given situation, and also can explain why they were retrieved and what factors came into play. Models related to decision-making and problem-solving in ACT-R over the past 10 years have seen increasing use of IBL to learn implicit knowledge structures. This is unsurprising given that ACT-R's declarative memory module and chunk structure is an excellent match for the storage and retrieval of instances, which effectively guides people to some form of IBL. In other words, the design and constraints of the architecture lead people to adopt an IBL-like approach by using the architecture in the most direct and intuitive way.

4. DISCUSSION

The goal of this data processing hierarchy is to increase the efficiency of data scientists by having each level of the hierarchy operate over the kinds of data on which they are most effective. Machine learning is massively parallel and operates best over large amounts of unstructured data to extract meaningful features and filter irrelevant data. The computing requirements are such that sequential extraction of hierarchical relations becomes impractical. Cognitive models require that their inputs have some structure over which their mechanisms can apply. While they have some of the parallelization characteristics of machine learning algorithms (in terms of matching large amounts of expert patterns to data), they also have the kinds of sequential and symbolic processing that can construct a richer knowledge structure with substantially less processing demands (e.g., Wigmorean evidence trees) [21].

We have recently developed a prototype cognitive model to categorize malware into families and extract intents [12] and compared it to a machine learning algorithm [19]. This model leverages the core features of ACT-R and uses an instance-based learning approach, with each sample of malware represented as a single instance in ACT-R's declarative memory. With one person working for less than one week of development, our model categorizes the full dataset in under 5 minutes and was able to exceed the machine learning algorithm's performance (accuracy 67.3%; recall precision, and F-1 score all exceeding 85.5%). That said, the basic malware families and intents (i.e., the features of each malware family) were provided to the model from human input. This kind of human ontological input is exactly the kind of complementary engagement that will prove successful in analyzing big data.

Another instance of the efficiency gained through that combination of statistical learning and symbolic (de)composition is our cognitive model of backgammon that is able to learn to perform at a highly skilled level after playing a few hundred games, as opposed to tens-of-thousands to millions of games for the equivalent machine learning algorithms to reach a comparable performance [18]. In general, we argue that cognitive models are a natural complement to machine learning techniques, providing scalable learning in structured information environments.

4.1 Shared Mental Models

Our argument that humans should best be used as supervisory control (and perform the high-level decision-making to which they are best suited) is not unique, and is similarly represented by other papers in this issue, such as [4]. For instance [12] examines the efficiency of human-in-the-loop analytics, and [5] argues for humans and machines/machine learning operating in continuous production loops (analogous to our *learning flows* and *data flows*).

The notion of shared mental models between humans and machines is a common thread when examining human-centered big data research. Mental models provide a representation of situation, various entities, capabilities, and past decisions/actions. These models are dynamic, with people and machines engaged in a continuous production loop [5][12]. Research on human teamwork in big data analysis [16] examines the degree to which teammates from different backgrounds have overlapping shared mental models, and the degree to which multiple agents can recognize a common plan from reading large corpora [3].

Cognitive models can provide a computational link between human and machine: they provide a quantitative, predictive understanding of shared mental models, and a cognitive computational basis for implementation of mental models in applied domains like robotics (see Figure 3), and as such support improved design of human-robot interaction tools and protocols.

ACT-R has recently been used to model a task that involved deciding on the allocation of responsibilities between a human and a robot teammate across several pursuit scenarios [10]. This model was designed to be able to parse any procedural information that may be described in a decision-tree form (i.e., an acyclic directed graph). That said, there is nothing about the control of this general-purpose model that precludes cyclic behavior, meaning that the model can theoretically perform any decision whose steps can be broken down into a directed graph.

The model takes a series of instructions, called *decision factors*, and accesses *factor values* for the current situation until it is instructed to make a decision. At that time, a decision chunk is retrieved either via standard retrievals with partial matching, or by blended retrievals through instance-based learning. The decision is also stored as a factor value for use in later decision-making steps. With the decision retrieved, the model moves onto the next decision in the chain.

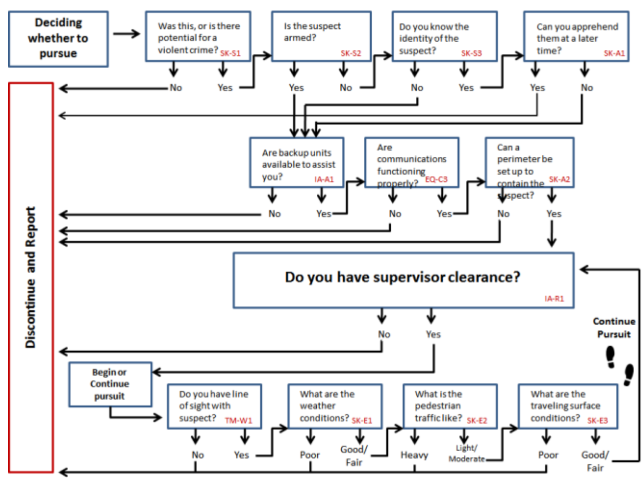

Figure 3. Example of Shared Mental Model from Robot Pursuit Task.

While the sequence of decision chains has to be provided as input to the model, they are derived from task instructions. The benefit of this general mental model is that it can process any arbitrary set of instructions using the same core productions. The model does not need to be changed in any way to tackle a new task but rather just needs a new set of instructions, just as human subjects do. In addition, it is possible to perform foraging behavior by having a pre-decision trigger that determines a value (such as expected information gain) that must reach a given threshold before moving onto the generation or revision of a decision.

To bring this discussion back to the particulars of big data, a cognitive model can then learn a human data analyst's search patterns and execute a shared plan to analyze big data. It can also take feedback from human operators to reinforce the utility of the learned procedures. This is analogous to the model exhibiting a naïve theory of mind of the human data scientist. A benefit of using cognitive models includes the offloading of data crunching from human data scientists onto the cognitive model. From some examples, we have shown that cognitive models are able to approximate the effectiveness and overall behavior of human in complex tasks without the same massive data requirements as traditional machine learning approaches. However, cognitive models still need further development to handle more automated abstraction of hierarchical data, and further research onto scalability will determine the best degree of abstraction for cognitive models to maximize their efficacy as a complementary tool in the analysis of big data.

5. SUMMARY

We argue that big data analysis is best performed with human oversight as a complementary piece of a hierarchy of cognitive and machine learning techniques that are computationally appropriate at their level of information complexity. Grounding intermediate knowledge structuring in a cognitive architecture provides for human-inspired reasoning with more flexible learning constraints (vis. parallelization) that can mitigate cognitive biases prevalent in comparable disciplines such as intelligence analysis. By doing so, we can increase the efficiency of the data analysis process by maximizing the effectiveness of structuring incoming data through massively parallel machine learning algorithms.

6. ACKNOWLEDGMENTS

This work is supported by the Intelligence Advanced Research Projects Activity (IARPA) via Department of the Interior (DOI) contract number D10PC20021. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. The views and conclusions contained hereon are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of IARPA, DOI, or the U.S. Government.

7. REFERENCES

[1] Ackoff, R. L. 1989. From data to wisdom. *Journal of Applied Systems Analysis*. 15, 3-9.

[2] Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., Qin, Y. 2004 An integrated theory of Mind. *Psychological Review, 111*, 1036-1060.

[3] Argenta, C., & Stewart, E. 2014. Extracting Short Stories from Large Data Sets. *HCBDR '14*.

[4] Campbell, S., Harbison, J. I., Bradley, P., & Saner, L. D. 2014. Cognitive Engineering Approaches to Analysis Training. *HCBDR '14*.

[5] Endert, A., & Pike, B. 2014. From Streaming Data to Streaming Insights: The Impact of Data Velocities on Mental Models. *HCBDR '14*.

[6] Gonzalez, C., Lerch, J. F., & Lebiere, C. 2003. Instance-based learning in dynamic decision making. *Cognitive Science, 27*, 591-635.

[7] Hellerstein, J. 2014. Data Science: From Hubris and Machismo to Human-Centered Design. http://www.trifacta.com/data-science-from-hubris-and-machismo-to-human-centered-design/ retrieved February 1, 2014.

[8] Klein, G., Moon, B., and Hoffman, R. R. 2006. Making sense of sensemaking 1: Alternative perspectives. *IEEE Intelligent Systems.* 21, 4, 70-73.

[9] Lebiere, C. 1999. The dynamics of cognition: An ACT-R model of cognitive arithmetic. *Kognitionswissenschaft.*, 8 (1), 5-19.

[10] Lebiere, C., Jentsch, F., & Ososky, S. 2013. Cognitive models of decision making processes for Human-Robot Interaction. *Proceedings of the HCI International Confernce,* July 2013.

[11] Lebiere, C., Pirolli, P., Thomson, R., Paik, J., Rutledge-Taylor, M., Staszewski, J., and Anderson, J. 2013. A Functional Model of Sensemaking in a Neurocognitive Architecture. *Computational Intelligence and Neuroscience; Special Issue on Neurocognitive Models of Sense Making.* DOI= http://dx.doi.org/10.1155/2013/921695.

[12] Liu, J., Wilson, A., & Gunning. D. 2014. Human-in-the-Loop Data Analytics System. *HCBDR '14.*

[13] Mandiant. 2013. APT1: Exposing One of China's Cyber Espionage Units. Mandiant Corp. URL: http://intelreport.mandiant.com/ retrieved January 21, 2014.

[14] Odom, W. E. 2008. Intelligence Analysis. *Intelligence and National Security.* 23, 3, 316-332.

[15] Oltramari, A. and Lebiere, C. (2012). Using Ontologies in a Cognitive-Grounded System: Automatic Action Recognition in Video Surveillance. *Proceedings of the 7th International Conference on Semantic Technology for Intelligence, Defense, and Security.*

[16] Paletz, S. B. F., 2014. Multidisciplinary Teamwork and Big Data. *HCBDR '14*.

[17] Rowley, J. 2007. The wisdom hierarchy: Representations of the DIKW hierarchy. *Journal of Information Science* 33, 163-180.

[18] Sanner, S., Anderson, J. R., Lebiere, C., and Lovett, M. 2000. Achieving efficient and cognitively plausible learning in backgammon. In *Proceedings of the Seventeenth International Conference on Machine Learning*. San Francisco: Morgan Kaufmann. 823-830.

[19] Shakarian, P. 2013. Cyber Adversarial Intent: A Frame-Based Logical Framework. *IARPA Technical Report.*

[20] Thomson, R. & Lebiere, C. 2013. Constraining Bayesian Inference with Cognitive Architectures: An Updated Associative Learning Mechanism in ACT-R. In *Proceedings of the 35th Annual Conference of the Cognitive Science Society*.

[21] Wray, R., Lebiere, C., Weinstein, P., Jha, K., Springer, J., Belding, T., Best, B., and Parunak, V. 2007. Towards a complete, multi-level cognitive architecture. In *Proceedings of the 8th International Conference on Cognitive Modeling.* Ann Arbor, MI.