Modeling Human Reasoning about Meta-Information

Sean Guarino, Jonathan Pfautz, & Zach Cox
Charles River Analytics Inc.
625 Mount Auburn Street
Cambridge, MA 02148

Emilie Roth
Roth Cognitive Engineering
89 Rawson Street
Brookline, MA 02445

Abstract

Information, as well as its qualifiers, or *meta-information*, forms the basis of human decision-making. Modeling human reasoning therefore requires the development of representations of both information and meta-information. However, while existing models and modeling approaches may include computational technologies that support meta-information analysis, they generally neglect its role in human reasoning. Herein, we describe the application of Bayesian Belief Networks to model how humans calculate, aggregate, and reason about meta-information when making decisions.

1. INTRODUCTION

Decision-making in real-time, dynamic environments is increasingly becoming a complex information management task, as new technologies generate ever-larger amounts of potentially relevant data. Decision-makers must therefore manage this incoming information, integrating it with previously gained knowledge to develop an understanding of the current situation (termed "situational awareness" by (Endsley & Garland, 2000; Endsley, 1995)). With this understanding, the decision-maker develops and selects a course of action that he or she believes will lead to a successful outcome. The ability to successfully decide on an effective course of action depends on the decision-maker's skill and experience in processing and understanding information. This ability fundamentally relies not only on understanding the domain-related information but also on the *qualities* of that information (e.g., recency, reliability), or the associated *meta-information*. Such qualities can critically influence how a decision-maker will process information, understand information, and choose a course of action based on that information.

Our analysis of cognitive tasks across different domains (e.g., wildfire management, military command and control, intelligence analysis, sensor management, weather impact analysis) has revealed that decision-makers reason using meta-information (Pfautz et al., 2006). As such, we have applied working definitions for terms that we use throughout this paper, as adapted from (Pfautz et al., 2005; Potter et al., 2000) (Potter, Roth, Woods and Elm, 2000):

- <u>*Data*</u> is output (processed or unprocessed) from a human or machine system that may or may not be useful in the decision-making process (e.g. radar reports atmospheric conditions of (x, y), Joe says a storm is coming, etc.)
- <u>*Information*</u> is recognized inputs that are necessary or usable in a directed decision-making process (e.g., a storm is coming that may affect the UAV's flight capabilities)
- <u>*Meta-data*</u> is characteristics or qualifiers of data (e.g., ground-based radar Y can only locate aircraft with an error of +/- 1.5m)
- <u>*Meta-information*</u> is characteristics or qualifiers of *information*, affecting a human's (or behavior model's) information processing, situation awareness, and decision-making (e.g., there is only 10% confidence that the threat is located at (x, y), therefore I will confirm its location before launching countermeasures)

These definitions are highly dependent on the particular cognitive task and the context in which that task is performed. Nevertheless, they serve to explicitly identify the critical role of meta-information in human decision-making. Our general approach to understanding the specific role of meta-information in this process involves an iterative application of *Cognitive Systems Engineering* (CSE) involving several phases of cognitive analysis, concept development, and user evaluation (Roth & Patterson, 2005; Potter, Gualtieri, & Elm, 2003; Potter et al., 2000).

Because of its critical role, it is essential to capture meta-information in advanced human behavior models such as SAMPLE (Harper et al., 2000), SOAR (Laird, 1987), or ACT-R (Anderson & Lebiere, 1998). While each of these models provide generic representations that will allow a savvy designer to integrate meta-information, none of them require or particularly encourage the inclusion of meta-information in models of human decision-making. In general, modelers do not address meta-information in these representations. Meta-information is not always available in the incoming data stream for these models, and may need to be separately obtained either through specific requests or additional computation. Once obtained, it would need to be integrated into a larger decision-making process (i.e. the role of track confidence in air combat threat analysis). In addition, types of meta-information are not always independent, meaning additional aggregation might be necessary before application to information processing, situation assessment, or decision-making processes. Clearly, incorporating meta-information in human behavior models represents a significant challenge.

In our efforts to model expert human decision-making behaviors using SAMPLE (Wallenstein et al., 2006; Harper et al., 2002; Harper et al., 2000), we have explored a number of approaches to the inclusion of meta-information, including rule-based behavior moderation (Aykroyd et al., 2005) and direct decision-making procedure modification (Guarino, Harper, & Zacharias, 2004). Often, in implementing decision-making processes described or demonstrated by subject matter experts (SMEs), we apply Bayesian Belief Networks (BBNs) to capture the situation assessment (SA) processes and meta-information gathering processes that individual decision-makers use to aggregate data and construct beliefs about their environment. In past efforts, we have applied meta-information to SA processes in several ways, including information filtering (e.g., determining which information behavior models should attend to), input calculations (e.g., moderating sensor readings based on meta-information about those sensors), and direct SA impact (e.g., additional nodes in SA models).

The focus of the research reviewed in this paper has been the exploration of Bayesian approaches to modeling reasoning about meta-information within our human behavior models. This research spans across several efforts and domains. In each of these efforts, one of our underlying goals was to understand the nature of the influence of meta-information, and generate approaches to modeling its impact. In Section 2, we describe relevant background material, including related material on human and computational reasoning about uncertainty. In Section 3, we cover methods for computing and aggregating meta-information from incoming data, as well as methods for incorporating meta-information into human behavior models. Finally, in Section 4, we present conclusions and directions for future work.

2. BACKGROUND

Most recent research into the kinds of difficulties presented by the need for decision-makers to reason about meta-information have been centered on uncertainty (Trickett et al., 2005; Sarter & Schroeder, 2001; Watkins, 2000; Bisantz, 1997; Teigen, 1988). We posit that uncertainty of information is only one type of qualifier that may affect information processing, situation awareness, and decision-making. Below, we discuss relevant research that has been focused on the role of uncertainty in human-decision-making and computational approaches to managing uncertainty. We also present our prior attempts to broadly define the types of meta-information we have encountered across different decision-making domains and a description of the SAMPLE architecture in which we develop human behavior models.

2.1 UNCERTAINTY AND HUMAN DECISION-MAKING

Human decision-making under uncertainty is recognized to deviate from classical, logical decision-making and to be based largely on experience-based heuristic methods (Kahneman, Slovic, & Tversky, 1982). Several attempts have been made to categorize different types of uncertainty and to identify how they affect the decision-making process. One method for classifying uncertainty is to look at its source (Booker, Anderson, & Meyer, 2003; Schunn, Kirschenbaum, & Trafton, 2003). Another method is to examine its use in the decision-making process, resulting in categories of uncertainty (Yovits & Abilock, 1974). Another set of classifications developed by Lipshitz and Strauss (1996) divides forms of uncertainty into inadequate understanding, lack of information, and conflicted alternatives. Similar taxonomies were developed by Schunn et al (2003) and Klein (1998). While these classifications of uncertainty and an understanding of their impacts on decision-making have been useful in the development of models of human behavior, they may not generalize to other types of meta-information not fundamentally based on uncertainty.

2.2 COMPUTATIONAL APPROACHES TO UNCERTAINTY

Computational systems have been developed to reason about uncertainties present in the real world in tasks ranging from weather forecasting to network security to financial risk management. To support this development, a variety of computational approaches have been developed to explicitly support reasoning about one or more types of uncertainty (Halpern, 2003; Parsons, 2001). These approaches include: probability measures, Dempster-Shafer belief functions (Russell & Norvig, 2003), extensions to first order logic (e.g., defeasible reasoning (McCarthy & Hayes, 1969), argumentation (Lin, 1993)), ranking functions, "plausibility" measures (Halpern, 2003), fuzzy set theory (Zadeh, 1965), and causal network methods (e.g., Bayesian belief networks (Pearl, 2001), similarity networks (Geiger & Heckerman, 1996), influence diagrams (Howard & Matheson, 1984)). This list, by no means exhaustive, represents the focus of computational research on the need to support automated reasoning about uncertainty.

Relatively recently, there has been increased interest in the management of *meta-data*, a term used to describe more broadly the various ways that data may be qualified (Havenstein, 2006; Marco & Jennings, 2004). This term has been applied to file systems, computer programs, images, relational databases, and data warehouses (i.e., its application is largely contained within the information technology community). Examples of meta-data include how, when, and by whom a particular set of data was collected, and how the data is formatted. This work has been focused on the tagging and handling of data according to its meta-data. Also, these efforts have been focused on the qualities inherent in the data rather than the qualities of the information that are used by a human decision-maker, and are thus less pertinent to our interest in modeling human behavior and meta-information.

2.3 SOURCES AND TYPES OF META-INFORMATION

In analysis efforts in previous work (Pfautz et al., 2005), we identified the main types of meta-information that impact the decision-making process (Figure 1). These types were derived from Cognitive Task Analysis (CTA) with a number of Subject Matter Experts (SMEs) in several military and commercial domains. We iteratively performed cognitive analysis, concept development, and user evaluation phases as part of a Cognitive Systems Engineering (CSE) approach with approximately 30 domain experts in over 400 hours of interviews. Based on a wide literature review (Burns & Hajdukiewicz, 2004; Eggleston, Roth, & Scott, 2003; Elm et al., 2003; Endsley, Bolte, & Jones, 2003; Vicente, 1999; Rasmussen, Pejtersen, & Goodstein, 1994), we believe this approach is necessary to gain a thorough understanding of the impact of meta-information in modeling human behavior.

The specific aspects of meta-information that are (or should be) considered by the decision-maker depend on the particular domain of application. By examining the *specific factors* that contribute to and constitute meta-information in the domains we examined, we were able to define a list of specific types of meta-information we encountered. While these types may or may not be applicable in other domains, they at least provide a useful aid in the identification of similar types of meta-information when analyzing different types of decision-making in other domains.

- Characteristics of the source of information
  - Type of data the source can produce
  - Frequency of reporting
  - Inherent biases in reporting
- Characteristics of the source as a function of other factors
  - Environmental effects (terrain, weather)
  - Information context
  - Time
- Uncertainty
  - Spatial uncertainties
  - Temporal uncertainties
  - Confidence
  - Uncertainties about uncertainty reporting
- Ambiguity
  - Specificity or resolution of information
  - Level of abstraction of information
- Information context (i.e., relationship to other information)
  - Degree of confirming or disconfirming information
  - Paucity of information
  - Missing or degraded information qualifiers
- Reliability of source
- Credibility of content from source
- Relevance or pertinence of information
- Temporal qualifiers
  - Staleness
  - Recency
  - Latency

Figure 1: Meta-Information Examples (Pfautz et al., 2006)

2.4 SAMPLE

SAMPLE (*Situation Assessment Model for Person-in-the-Loop Evaluation*) is a domain-independent architecture for modeling situation awareness (SA) centered decision-making in high-stress, time-critical environments. SAMPLE is a general-use human behavior model that has been recently applied to the commercial aviation arena under a number of efforts in the Distributed Air/Ground Traffic Management (DAG-TM) domain (Harper et al., 2002) and to Military Operations on Urban Terrain (MOUT) (Aykroyd et al., 2002).

A SAMPLE agent draws information from the world (real or simulated), filters that information according to the agent's current attentional focus, and processes the information to generate a set of identified events. A situation assessment process then translates the low-level events into high-level situation assessments, which are finally processed by a decision-making module to produce actions or communications. Actions chosen by the decision-making module are passed to a goal manager/action selector, which deconflicts and prioritizes the agent's currently chosen tasks. Actions taken by the agent that redirect attentional focus are sent back to the attention allocator module, creating a feedback loop.

Each of the component cognitive modules draws from a suite of internalized mental models of the external world stored in long-term memory. Additionally, each module both draws from and populates a short-term memory representation which collectively model the individual's real-time interpretation of the world state. Cognitive processes are modeled computationally through several AI technologies, including Bayesian Belief Networks.

3. BAYESIAN APPROACHES TO MODELING META-INFORMATION

In this effort, we analyzed a range of approaches we have taken to modeling meta-information in previous efforts. In exploring these methods, we drew upon extensive past research in which we modeled human behavior and indirectly included the impact of meta-information upon that behavior (Pfautz et al., 2005; Aykroyd et al., 2005; Guarino et al., 2004; Mulgund et al., 1996). Herein, we review these approaches, and provide analysis on their capability to address problems associated with modeling meta-information. In each of the efforts that generated these examples, meta-information modeling was not a central goal, and its consideration in the cognitive modeling process was limited to the ways in which domain experts described or demonstrated the impacts of meta-information on their decision-making.

3.1 TYPES OF REASONING TO MODEL

The first issue we encountered in our exploration of the application of BBNs to meta-information analysis was concerned with the type of modeling support needed. Given that many techniques inherently provide some model of a phenomena or phenomenon, we needed to identify the elements in the domain to be captured in a BBN. In the problem domain, there are a number of different types of reasoning that can be captured, from deductive reasoning about factors predicting an event to abductive reasoning about the degree of support for a particular hypothesis. BBNs support both types of reasoning, and therefore can be used to model abduction, deduction, or both. In a practical sense, this means we can generate BBNs to support modeling of the expert's reasoning (e.g., deductive reasoning about the expert's confidence in a sensor given meta-information surrounding the sensor report, as shown in Figure 2a) or modeling of the sensor's error likelihood (e.g., reasoning abductively about sensor error from various factors affected by that error, as shown in Figure 2b).

To apply BBNs in either of these manners, one must have an understanding of not only the domain, but also the ramification of applying different modeling techniques to the problem. In a generic sense, the application of a particular type of reasoning to a problem may seem trivial; e.g., reasoning about the support for hypothesis given some evidence vs. reasoning about the likely outcome given some evidence. However, the nature of many problem domains is such that multiple types of reasoning could be used effectively for the same problem and only the semantics of the application will differ. This problem requires further investigation as additional computational techniques for supporting met+a-information analysis are explored.

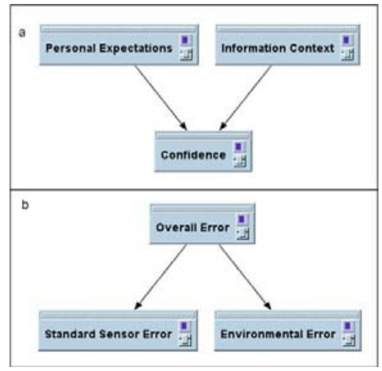

Figure 2: Deductive vs. Abductive Reasoning: a) Deductive Reasoning; b) Abductive Reasoning

3.2 CALCULATING META-INFORMATION FROM DATA AND META-DATA

Meta-information can be obtained in a number of ways by a decision-maker. In many cases, a human decision-maker will have to compute meta-information from multiple factors (e.g., in a poker game, is Bob bluffing if his eye twitches and he shifts frequently when he raises?). While some systems have the ability to produce meta-data about their performance (e.g., a tool of type X has an error of ± 0.5), only in particular tasks can that meta-data be used directly as meta-information. Generally, however, humans develop meta-information during their reasoning process. Any cognitive model must behave similarly.

BBNs can be directly applied to model this cognitive computation of meta-information, as illustrated in Figure 3. This BBN does not explicitly include the effect of the sensor type, or other moderating information, on these environmental factors. Moderators can be added in several different ways, including the use of separate BBNs based on different systems (for example, perhaps a second sensor is unaffected by lighting conditions, as illustrated in Figure 4), modifications to the conditional probability tables (CPTs) within the network (for example, perhaps a second sensor is more sensitive to weather changes), or, in more simplistic cases, the addition of a moderator as an additional factor in the network. These cognitive computations often must be performed over space and time (e.g., How well does the sensor perform in this region, with these terrain restrictions? How does performance degrade as the battery runs down?), creating a need for enhanced modeling methods such as Dynamic Bayesian Belief Nets (DBNs).

Because meta-information types are not inherently independent (e.g., the type of information being considered interacts with the type of source providing the information to influence reliability), meta-information may need to be aggregated and combined. One approach to combining meta-information of various types is aggregation through the application of BBNs. Figure 6 provides an example of aggregating different types of meta-information, where the effect of information on confidence (e.g., is there lots of supporting or conflicting information from other data?), the effect of expectations on confidence (e.g., does the given information fit what the behavior model expects to occur?), and the overall sensor error (e.g., combination of sensor obstructions, including weather, terrain, etc.) are combined into an overall confidence that impacts the usefulness of specific sensor reports.

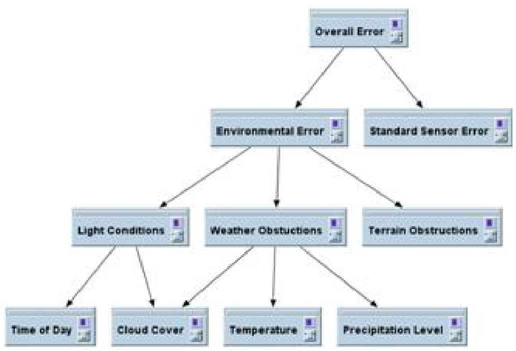

Figure 3: Computing Environmental Error Meta-Information From Environmental Data

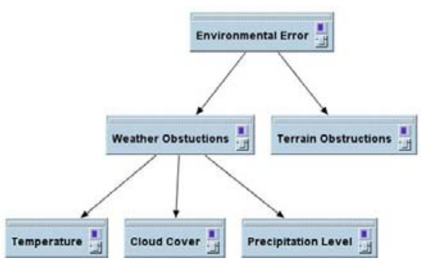

Figure 4: Environmental Error for Sensor Unaffected by Lighting

Like the computation of a specific type of meta-information, knowing the best means to aggregate meta-information is challenging. Observation and study of human decision-making may provide some justification, but could result in inadvertent inclusion of biases (e.g., predisposition to particular source types). On the other hand, using engineering data about sources may not adequately represent how a human would reason about meta-information. Therefore, the development of methods that model how decision-makers aggregate meta-information is an open research challenge that we continue to investigate. Furthermore, the influence of meta-information on how a behavioral decision should be made (or how a piece of information should be assessed) needs to be further investigated to understand what types of meta-information would normally be calculated by humans and, particularly, the impact of *not* calculating meta-information on the accuracy of the decision-making process.

Figure 5: Sensor Type As Node in Network: a) Sensor 3 More Sensitive to Terrain Obstruction; b) Sensor 1 More Sensitive to Weather Obstruction

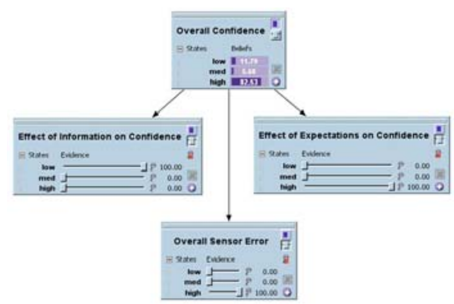

Figure 6: Aggregating Meta-Information to Compute Overall Confidence

3.3 APPLYING META-INFORMATION IN COGNITIVE MODELS

If we assume that meta-information must be computed and/or aggregated in some tractable manner by the decision-maker, then the next step is to understand how meta-information could be used in modeling a reasoning process. One approach is to simply apply meta-information in ways that filter or prioritize information based on meta-information. For example, when receiving a large number of incoming sensor reports, we might limit the reports impacting a cognitive model based on results of meta-information analysis. This approach requires some degree of cognitive task analysis (and/or human-in-the-loop experimentation) to determine how a subject matter expert would perform this filtering or prioritization based on the given meta-information, as well as the current decision-making task and the current situation. However, because attentional filtering mechanisms have already been modeled in some detail (Harper & Zacharias, 2002), extending them to incorporate filtering/prioritization based on meta-information is a relatively simple application of the previously discussed meta-information calculations.

Another way to incorporate meta-information is to include it within the BBN models of information gathering, situation assessment, and decision-making processes. This involves generating BBNs to base actions on the application of meta-information (e.g., if threat report is sufficiently recent, then act on the report) and/or changing internal data representations according to meta-information (e.g., if confidence is high, then interpret data with more precision). For example, a BBN used by air-combat pilots to analyze the location of an approaching track might be enhanced with a meta-information node aggregating the track confidence, as illustrated in Figure 7. In this example, when information is posted informing the model that the confidence is low, the threat level increases, representing the human likelihood of worrying more about threats for which there is little information available. This will result in the model emphasizing information gathering for this particular track, which may lead to an increase in track confidence and, ultimately, a more accurate calculation of the threat posed by the track.

This approach allows the (fuzzified) meta-information to be incorporated into the cognitive reasoning process and allows some explicit control over its influence. Furthermore, with DBNs, it could incorporate the influence over time and handle multiple types of influence (e.g., inhibitory, excitatory) on other variables. However, because of the number of potential types of meta-information, this approach may rapidly overload the representation of the BBN, increasing its computational complexity and obfuscating its purpose. Another approach is to use the meta-information in a specific parameter; in BBNs, this means directly changing the probability on a node as a function of meta-information. Examples of how the probability of evidence could be alternatively computed are shown in Table 1. Here, each of the calculated probabilities represents one possible discrete value for a node in a BBN. Rather than setting a information values value in behavior modeling BBNs based purely on incoming sensor information, the information is modified externally based on meta-information such as sensor reliability, information confidence, credibility, and sensor type. For example, in Table 1, the probability that the location of a detected entity is "near" is some sensor value (K) multiplied by a meta-information representation of the *reliability* of that sensor. Lethality of an enemy unit might be similarly calculated based on some human intelligence report (J) multiplied by the behavior model's *confidence* in that human intelligence. These computational formulas for calculating BBN values can be as complex as necessary, potentially being represented through some function of sensor value, sensor type, and a host of available meta-information concerning that sensor and related environmental information.

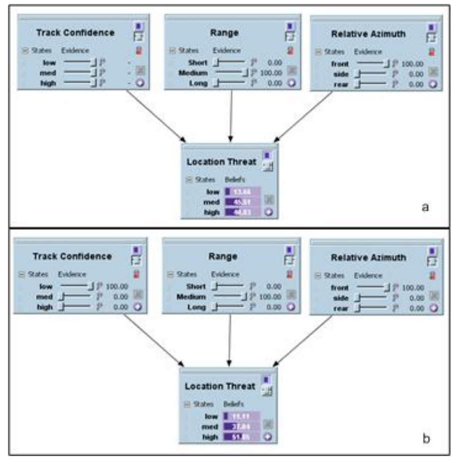

Figure 7: Incorporating meta-information explicitly into a BBN: a) no confidence information posted; b) low confidence posting increases analyzed threat

Table 1: Examples of Computing the Probability of a Discrete Value for a BBN Node as a Function of Associated Meta-Information

Probability (Location = "near")	=	K * Reliability
Probability (Lethality = "low")	=	J* Confidence
Probability (Threat = "high")	=	f (value, type, credibility)

This approach, in which meta-information is managed externally from the BBN (and not explicitly captured in the graphical representation) can reduce some computational complexity. However, hiding the intermediate calculations through which meta-information is integrated could obfuscate the representation of the reasoning process, and substantially limit the robustness with which meta-information is integrated into the behavior model.

4. CONCLUSIONS

A key aspect of modeling human behavior is capturing the effect of meta-information on information processing, situation assessment, and decision-making. Our experience performing cognitive task analyses in different decision-making domains has shown that experts repeatedly use this meta-information when making decisions. Although many advanced human behavior models have methods by which meta-information could be explicitly represented (e.g., rules in SOAR or ACT-R, BBNs in SAMPLE, etc.), none of these models require or even encourage the inclusion of meta-information when modeling human behavior. In this effort, we have begun to explore approaches to modeling meta-information generation and application using BBNs. Each of these approaches has been applied within SAMPLE agents to more accurately model decision-making processes. We described the application of BBNs in modeling each of the following types of cognitive tasks:

- Computation of complex meta-information through aggregation of data and meta-data
- Computation of meta-information through aggregation of other types of meta-information
- Application of meta-information to BBNs representing information gathering, situation assessment, and decision-making processes
- Application of meta-information to modify information values before applying them in required cognitive processes

These approaches clearly indicate how BBNs can provide an effective tool for modeling and application of meta-information in cognitive modeling efforts.

In addition, this research has indicated a need to more carefully include the influence of meta-information when designing complex human behavior models. In future efforts, we foresee the application of these approaches within our own SAMPLE agents, and recommend the inclusion of meta-information within the wide range of cognitive models applied within other modeling architectures.

Acknowledgements

The work described in this paper was performed, in some part, under OSD Contract No. FA8650-04-M-6418, U.S. Army Contract No. DAAB07-02-C-L403, OSD Contract No. FA8650-04-D-6549, OSD Contract No. W911QX-04-C-0063, and NASA Contract No. NNS04AA26C. The authors would like to express their deepest gratitude to the different domain experts interviewed and observed as part of their analysis. While many experts must necessarily remain anonymous, we would like to thank Ted Fichtl for his contributions.

References

Anderson, J.R. & Lebiere, C. (1998). *Atomic Components of Thought*. Hillsdale, NJ: Erlbaum Associates.

Aykroyd, P., Bachman, J., Harper, K., & Hudlicka, E. (2005). Modeling Individual Differences and Stressors Using the SAMPLE Cognitive Architecture. In *Proc. of 14th Conf. on Behavior Representations in Modeling and Simulation. Universal City, CA.*

Aykroyd, P., Harper, K.A., Middleton, V., & Hennon, C.G. (2002). Cognitive Modeling of Individual Combatant and Small Unit Decision-Making Within the Integrated Unit Simulation System. In *Proc. of 11th Conf. on Computer-Generated Forces and Behavior Representation (CGF-BR),* (pp. 597-604). Orlando, FL.

Bisantz, A.M. (1997). *Modeling Environmental Uncertainty to Understand and Support Dynamic Decision Making.* Georgia Int. of Technology, School of Industrial and Systems Engineering, Atlanta, GA.

Booker, J.M., Anderson, M.C., & Meyer, M.A. (2003). The Role of Expert Knowledge in Uncertainty Quantification (Are We Adding More Uncertainty or More Understanding?). In *Proc. of Seventh U.S. Army Conf. on Applied Statistics*.

Burns, C.M. & Hajdukiewicz, J. (2004). *Ecological Interface Design*. CRC Press.

Eggleston, R., Roth, E., & Scott, R. (2003). A Framework for Work-Centered Product Evaluation. In *Proc. of Human Factors and Ergonomics Soc. 47th Ann/ Meeting*. Denver, CO: Human Factors and Ergonomics Soc.

Elm, W., Potter, S., Gualtieri, J., Roth, E., & Easter, J. (2003). Applied CWA: A Pragmatic Methodology for Designing Revolutionary Cognitive Affordances. In E. Hollnagel (Ed.), *Handbook for Cognitive Task Design*. London: Lawrence Erlbaum Associates.

Endsley, M.R. (1995). Measurement of SA in Dynamic Systems. *Human Factors*, 37(1), 65-84.

Endsley, M. R., Bolte, B., & Jones, D. G. (2003). *Designing for Situation Awareness: An Approach to Human-Centered Design*. London: Taylor & Francis.

Endsley, M. R. & Garland, D. J. (2000). *Situation Awareness, Analysis, and Measurement*. Mahwah, NJ: Lawrence Erlbaum Associates.

Geiger, D. & Heckerman, D. (1996). Knowledge Representation and Inference in Similarity Networks and Bayesian Multinets. *Artificial Intelligence*, 82(45), 74.

Guarino, S. L., Harper, K. A., & Zacharias, G. L. (2004). *Distributed Air/Ground Negotiation Optimization Final Report.* (Rep. No. R01664). Charles River Analytics Inc.

Halpern, J. (2003). *Reasoning About Uncertainty*. Cambridge, MA: MIT Press.

Harper, K., Guarino, S., White, A., Hanson, M. L., Bilimoria, K., & Mulfinger, D. (2002). An Agent-Based Approach to Aircraft Conflict Resolution With Constraints. In *Proceedings of AIAA Guidance, Navigation, and Control*. Monterey, CA.

Harper, K. A., Mulgund, S. S., Zacharias, G. L., & Menke, T. (2000). SAMPLE: Situation Awareness Model for Pilot-in-the-Loop Evaluation. In *Proceedings of 9th Conference on Computer Generated Forces and Behavioral Representation*. Orlando, FL.

Harper, K.A. & Zacharias, G.L. (2002). Modeling Attention Allocation and Multitasking in Computational Human Behavior Representations. In *Proc. of 11th Conf. on Computer Generated Forces and Behavior Representation*. Orlando, FL.

Havenstein, H. (2006, July 18). Metadata Management Returns to the Fore. *ComputerWorld*.

Howard, R. & Matheson, J. (1984). Influence Diagrams. In R. Howard & J. Matheson (Eds.), *Readings on the Principles and Applications For Decision Analysis* (pp. 719-762). Menlo Park, CA: Strategic Decisions Group.

Kahneman, D., Slovic, P., & Tversky, A. (1982). *Judgment Under Uncertainty: Heuristics and Biases*. Cambridge, UK: Cambridge University Press.

Klein, G. A. (1998). *Sources of Power: How People Make Decisions*. Cambridge, MA: MIT Press.

Laird, J. E. (1987). Soar: An Architecture for General Intelligence. *Artificial Intelligence,* 33(1), 1-64.

Lin, F. (1993). An Argument-Based Approach to Non-Monotonic Reasoning. *Comp. Intelligence,* 9254-267.

Lipshitz, R. & Strauss, O. (1996). How Decision-Makers Cope With Uncertainty. In *Proc. of the Human Factors and Ergonomics Society 40th Annual Meeting-1996,* (pp. 189-193). Philadelphia: Human Factors Society.

Marco, D. & Jennings, M. (2004). *Universal Meta-Data Models*. New York, NY: Wiley.

McCarthy, J. & Hayes, P.J. (1969). Some Philosophical Problems From the Standpoint of Artificial Intelligence. In B. Meltzer, D. Mitchie, & M. Swann (Eds.), *Machine Intelligence 4* (pp. 463-502). Edinburgh, Scotland: Edinburgh University Press.

Mulgund, S., Rinkus, G., Illgen, C. & Zacharias, G. (1996). *Situation Driven Workload Adaptive Pilot/Vehicle Interface.* (Rep. No. R96051). Charles River Analytics, Cambridge, MA.

Parsons, S. (2001). *Qualitative Methods for Reasoning Under Uncertainty*. Cambridge, MA: MIT Press.

Pearl, J. (2001). *Causality: Models, Reasoning, and Inference*. Cambridge Univ Press.

Pfautz, J., Roth, E., Bisantz, A., Fouse, A., Madden, S., & Fichtl, T. (2005). The Impact of Meta-Information on Decision-Making in Intelligence Operations. In *Proceedings of Human Factors and Ergonomics Society Annual Meeting*. Orlando, FL.

Pfautz, J., Roth, E., Bisantz, A., Llinas, J., & Fouse, A. (2006). The Role of Meta-Information in C2 Decision-Support Systems. In *Proc. of Command and Control Research and Technology Symposium*. San Diego, CA.

Potter, S. S., Gualtieri, J. W., & Elm, W. C. (2003). Case Studies: Applied Cognitive Work Analysis in the Design of Innovative Decision Support. In E. Hollnagel (Ed.), *Handbook for Cognitive Task Design*. London: Lawrence Erlbaum Associates.

Potter, S. S., Roth, E. M., Woods, D. D., & Elm, W. C. (2000). Bootstrapping Multiple Converging Cognitive Task Analysis Techniques for System Design. In Chipman, Shalin, & Schraagen (Eds.), *Cognitive Task Analysis* (pp. 317-340). New Jersey: Lawrence Erlbaum.

Rasmussen, J., Pejtersen, A., & Goodstein, L. (1994). *Cognitive Systems Engeering*. NYC, NY: Wiley and Sons.

Roth, E. M. & Patterson, E. S. (2005). Using Observational Study As a Tool for Discovery: Uncovering Cognitive and Collaborative Demands and Adaptive Strategies. In H. Montgomery, R. Lipshitz, & B. Brehmer (Eds.), *How Professionals Make Decisions*. Mahweh, NJ: Lawrence Erlbaum Associates.

Russell, P. & Norvig, P. (2003). *Artificial Intelligence: A Modern Approach*. (Second ed.) Upper Saddle River, NJ: Prentice-Hall.

Sarter, N. B. & Schroeder, B. (2001). Supporting Decision Making and Action Selection Under Time Pressure and Uncertainty: The Case of in-Flight Icing. *Human Factors,* 43573-583.

Schunn, C. D., Kirschenbaum, S. S., & Trafton, J. G. (2003). The Ecology of Uncertainty: Sources, Indicators, and Strategies for Information Uncertainty.

Teigen, K. H. (1988). The Language of Uncertainty. *Acta Psychologica,* 6827-38.

Trickett, S., Trafton, G., Saner, L., & Schunn, C. (2005). "I Don't Know What Is Going on There": The Use of Spatial Transformation to Deal With and Resolve Uncertainty in Complex Visualizations. In M. Lovett & P. Shah (Eds.), *Thinking With Data*. Mahwah, NJ: Erlbaum.

Vicente, K. J. (1999). *Cognitive Work Analysis: Towards Safe, Productive, and Healthy Computer-Based Work.* Mahwah, NJ: Lawrence Erlbaum Associates.

Wallenstein, C., Reilly, S. N., Harper, K. A., & McNally, B. (2006). Building a Human Behavior Model for a Collaborative Air-Combat Domain. In *Proc. of BRIMS*.

Watkins, E. T. (2000). *Improving the Analyst and Decision-Maker's Perspective Through Uncertainty Visualization.* Master of Science, Air Force Institute of Technology, Dept. of Systems and Engineering Management.

Yovits, M. C. & Abilock, J. (1974). A Semiotic Framework for Informaiton Science Leading to the Development of a Quantitative Measure of Information. In *Proc. of 37th American Society for Information Sciences Meeting,* (pp. 163-168).

Zadeh, L. A. (1965). Fuzzy Sets. *Information and Control,* 8338-353.