International Journal of Approximate Reasoning 50 (2009) 437–449

Contents lists available at ScienceDirect

International Journal of Approximate Reasoning

journal homepage: www.elsevier.com/locate/ijar

Modeling human reasoning about meta-information

Sean L. Guarino [a,*], Jonathan D. Pfautz [a], Zach Cox [a], Emilie Roth [b]

[a] Charles River Analytics, 625 Mount Auburn Street, Cambridge, MA 02138, United States
[b] Roth Cognitive Engineering, 89 Rawson Street, Brookline, MA 02445, United States

ARTICLE INFO

Article history:
Received 1 February 2007
Received in revised form 14 September 2007
Accepted 16 April 2008
Available online 17 June 2008

Keywords:
Human behavior representations
Meta-information
Bayesian belief networks

ABSTRACT

Information, as well as its qualifiers, or *meta-information*, forms the basis of human decision-making. Human behavior models (HBMs) therefore require 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.

Published by Elsevier Inc.

1. Introduction

Human 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 (sometimes termed "situational awareness" [1,2]). 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 qualifiers, or associated *meta-information*, describing that information (e.g., recency, reliability, source, etc.). Such qualities contextualize information, and therefore can critically influence how a decision-maker will process, understand, and act on that information. For example, Suzy decides to attend the new contemporary art museum despite an email from Rob describing it as "boring" because this comment comes from Rob, who dislikes modern art. The information "the museum is boring" is qualified by its source, Rob, and Suzy's reasoning is impacted by her prior knowledge of that source. If the source of the information changes, or knowledge about that sources changes, the information may result in different perceptions, reasoning, and action from Suzy.

This simple example represents but one of many cases where we have explored the role of meta-information in human reasoning. Our analysis of cognitive tasks across different domains (e.g., wildfire management, military command and control, intelligence analysis, sensor management, weather impact analysis, among others) has revealed that decision-makers reason using meta-information [3]. This research is substantiated by the considerable literature on reasoning under uncertainty, which we consider to be one form of meta-information [4–6]. As such, we have developed working definitions for terms that we use throughout this paper, as adapted from [7,8]

* Corresponding author. Tel.: +1 617 491 3474x561; fax: +212 617 868 0780.
E-mail addresses: sguarino@cra.com (S.L. Guarino), jpfautz@cra.com (J.D. Pfautz), zcox@cra.com (Z. Cox).

0888-613X/$ - see front matter Published by Elsevier Inc.
doi:10.1016/j.ijar.2008.04.009

- *Data* 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, Joe says a storm is coming, etc.).
- *Information* is recognized inputs that are necessary or usable in a directed decision-making process or behavior (e.g., a storm is coming that may affect the UAV's flight capabilities).
- *Meta-data* is characteristics or qualifiers of *data* that may or may not be useful in the decision-making process (e.g., ground-based radar Y can only locate aircraft with an error of +/−1.5 m).
- *Meta-information* is characteristics or qualifiers of *information*, affecting a human's (or a model of a human's) decision-making, reasoning, or behavior:
  - Information processing (e.g., reports flagged as "important" are used first).
  - Situation awareness (e.g., because information about wind speed is recent and certain, the model can ascertain which towns are threatened by tornados).
  - Decision-making (e.g., because information about the adversary's location is 30 h old, the model must actively gather new information before moving into that location).

Categorizing inputs according to these definitions is dependent on the particular cognitive task and the context in which that task is performed. Nevertheless, the definitions 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), a methodology for defining aspects of human reasoning and behavior to aid system design that involves several phases of cognitive analysis, concept development, and user evaluation [9–11].

Because of its role in reasoning, and particularly in human reasoning, any agent or *human behavior model* (HBM) that attempts to replicate human cognitive processes (e.g., advanced models such as SAMPLE [12], SOAR [13], or ACT-R [14]) must necessarily capture the impacts of meta-information on those processes. HBMs are useful in a wide variety of applications, including both theoretical (e.g., developing and testing theories of human cognition, including emotions, perception, decision-making, and action performance) [15–18] and practical (e.g., representing realistic human behavior in training and other simulation environments, such as games, tracking human behavior to automatically adapt decision-support, or simply managing complex tasks normally performed by humans) [19–21].

These applications will necessarily span domains where it is critical to incorporate models of how humans reason about meta-information. While each of the existing HBMs mentioned above provides generic representations that will allow a savvy designer to integrate meta-information, none of them require or particularly encourage the inclusion of meta-information. Furthermore, modelers typically 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 [19,22,23], we have explored a number of approaches to the inclusion of meta-information, including rule-based behavior moderation and direct decision-making procedure modification [24]. 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 [19,24], 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 [25–29]. We posit that uncertainty of information is only one type of qualifier that may affect information processing, situation awareness (or understanding), 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 [30]. Often, these heuristic approaches represent how experts reason about the meta-information surrounding these decision-making processes, and are therefore crucial to capture in accurate decision-making models. 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, for instance, dividing uncertainty into forms that come from computational models as opposed to human interpretation [31]. Another method is to examine its use in the decision-making process, resulting in categories of uncertainty, which has resulted in categories such as *executional uncertainty, goal uncertainty*, and *environmental uncertainty* [32]. Another set of classifications developed by Lipshitz and Strauss [33] divides forms of uncertainty into *inadequate understanding*, *lack of information*, and *conflicted alternatives.* Similar taxonomies were developed by Schunn et al. [31] and Klein [34]. These taxonomies can prove to be useful in attempts to develop descriptive models of human reasoning. For example, Lipshitz and Straus [33] discuss five strategies for reasoning under uncertainty: (1) reduce uncertainty by collecting more information; (2) use assumptions to fill in gaps of knowledge; (3) weigh pros and cons; (4) forestall; and (5) suppress uncertain information. 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 (e.g., factors such as pedigree or recency which also serve to contextualize information and impact how it is perceived, understood, and used to drive behavior).

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 [35,36]. These approaches include: probability measures, Dempster–Shafer belief functions [37], extensions to first-order logic (e.g., defeasible reasoning [38], argumentation [39]), ranking functions, "plausibility" measures [35], fuzzy set theory [40], and causal network methods (e.g., Bayesian belief networks [41], similarity networks [42], influence diagrams [43]). Within these approaches, additional methods have been developed for aggregating and propagating uncertainty (e.g., computing "second-order uncertainty" in Bayesian networks) [44]. This list, by no means exhaustive, represents the focus of computational research on the need to support automated reasoning about uncertainty [45].

Some effort has been made, as part of the development of these approaches, to define uncertainty and to describe taxonomies of uncertainties that computational systems may reason about. Of these taxonomies, Smets [46], Smithson [47], and Bosc and Prade [48] are notable. However, examination of these (and other attempts to structure the meaning of "uncertainty"), and the large variation in definitions and taxonomies supports Elkan's [49] assertion that developing such taxonomies is largely a philosophical exercise. This assertion may relate to the degree to which the development of these computational approaches are tied to an understanding of (and desire to model) human reasoning, particularly relative to a particular task or context. Within a specific task (or class of tasks), it may be possible to study the impacts of uncertainty, and meta-information more generally) on human reasoning, and then apply these techniques to model that reasoning.

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 [50,51]. 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 (e.g., a typical email header contains many examples of meta-data). This work has been focused on the tagging and handling of data according to its meta-data with little linkage to human reasoning about that meta-data (e.g., how do the components of the email header contribute to the order in which a human might process their email?). Because 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 to process, understand, and act, they are less pertinent to our interest in modeling how meta-information may impact human reasoning and behavior.

2.3. Sources and types of meta-information

In analysis efforts described in previous work [8], we identified the main types of meta-information that impact the decision-making process within a set of well-defined application domains (Table 1). These types were derived from cognitive task analysis (CTA) [11,52] with a number of subject matter experts (SMEs), a systematic process that reveals the characteristics of the work domain as well as the cognitive process used to perform work through observation and interviewing techniques. Across these analysis efforts, we also developed prototype concepts for decision-support, and conducted user

Table 1
Sources and types of meta-information in explored domains (see Pfautz et al. [3] for a more detailed analysis)

Meta-information type	Sub-types or related types	Example impacts on decision-making
Characteristics of the information source	Type of data the source can produce Type of processing used Range of data generated Baseline error rates Frequency of reporting Ability to report its status and characteristics of that report Inherent biases Past performance and history Directly observing or deriving information	Because Pete is extremely experienced, and has been largely successful in the past, Pete's reports on new business directions will be given more weight by senior management
Characteristics of the source as a function of other factors	Time Location in environment (e.g., terrain and weather) Types of intermediate processing Content of report	Because his superiors know that Jim is located in the middle of a major dust storm, he is likely to miss observing a passing enemy, leading his superiors to disregard his report that nobody has passed
Uncertainty	Spatial uncertainties Temporal uncertainties Uncertainties about uncertainty reporting Likelihood Probability Confidence Accuracy Precision	Emma's estimate of the number of individuals in a large crowd is imprecise because the crowd was constantly changing in size over time as people came and went, therefore the event organizers prepared extra food
Ambiguity	Specificity or resolution of information Level of abstraction of information	Juan's report that a car is coming up the road is not enough for Bobby to start making dinner, because it's not clear if the car is actually the car belonging to their dinner guests
Information context (i.e., relationship to other information)	Degree of confirming or disconfirming information Paucity of information Frequency of reporting of information Missing or degraded information qualifiers Information-to-noise ratio	Without any information on the quality of her water, Jane would not drink it. Because both Rakesh and Bill independently tested and confirmed the water quality, Jane feels safe drinking it
Reliability of source	W.r.t. source characteristics W.r.t. information context	Channel 5's weatherperson is the most experienced at predicting snow, so Charles always tunes into Channel 5 on cold days before he decides to drive his sports car or pick-up truck
Credibility of content from source	W.r.t. reliability W.r.t. type of content W.r.t. type of source W.r.t. information context	Chen never uses Bob's recommendations on which football team will win in his office pool, because he knows Bob does not know anything about football
Relevance or pertinence	W.r.t. specific mission goals W.r.t. actual/perceived information needs W.r.t. broader operational context W.r.t. current hypotheses about the situation	Sam's reports about the weather have no impact on the decision by the schoolkids about whether to play volleyball or kickball because they are playing inside in the gymnasium
Temporal qualifiers	Staleness Recency Certainty about time of reporting Latency Lag	While Ed may be trustworthy, his report from 6 h ago about, where the weather balloon is may not be accepted, as the balloon has likely drifted a significant distance in that time

evaluation to further refine our understanding of the cognitive mechanisms in use in the work domain. This research encompassed interactions with over 30 domain experts (in the different domains) and over 500 h of interviews, observation, and evaluation. Based on our experience in human behavior modeling, knowledge elicitation, and the supporting literature [53–57], we believe that this overall approach to developing an understanding of expert reasoning and behavior is sufficient to begin to understand the impact of meta-information in human cognition at a level that supports modeling. One product of these analyses has been the development of a list of sources and types of meta-information we have consistently encountered across application domains. As opposed to the taxonomies of uncertainty discussed earlier, these sources and types come strictly from the study of human reasoning.

In our analysis, we discovered that the specific aspects of meta-information that are (or should be) considered by the decision-maker depend on the particular domain of application, and the particular task being performed in that domain or 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, and provide the basis for beginning to model the impacts of meta-information on human reasoning.

2.4. Modeling human reasoning and behavior

The computational representation of human reasoning and behavior has applications in a range of domains, supporting training, modeling and simulation, as well as efforts simply to better understand human perception, reasoning, and action. Many approaches to modeling have been developed and documented (see Pew and Mavor [58] for an overview of a subset of these approaches). In this paper, we discuss human behavior modeling with respect to our own efforts to model human cognition and behavior based on recognition-primed decision-making [59]. Fig. 1 illustrates our modeling approach, SAMPLE (*situation assessment model for person-in-the-loop evaluation*), which is a domain-independent architecture developed for modeling situation awareness-centered decision-making in high-stress, time-critical environments, based largely on Klein's theory of recognition-primed decision-making [60]. Recognition-primed decision-making posits that experts do not do significant amounts of reasoning and problem solving, but rather have been trained to recognize the criticial elements of a situation and to act accordingly. Functionally, SAMPLE is a general-use HBM that has been applied in a variety of domains, including the commercial aviation arena in air/ground traffic management simulations [23], in the modeling of adversary pilots in military simulations [12], and in the modeling of the impacts of stressors and individual differences in small unit military operations in urban terrain.

SAMPLE communicates with a simulated environment (the "World") through sensors and actions. Inputs to SAMPLE are first processed by an *information processing* module. This module is typically built using Fuzzy logic [61] components that turn the real-valued sensor data into fuzzy membership values that are more compatible with the way human decision-makers tend to reason about domains. In more complex situations, it can include a filtering system capable of simulating the agent's current attentional focus.

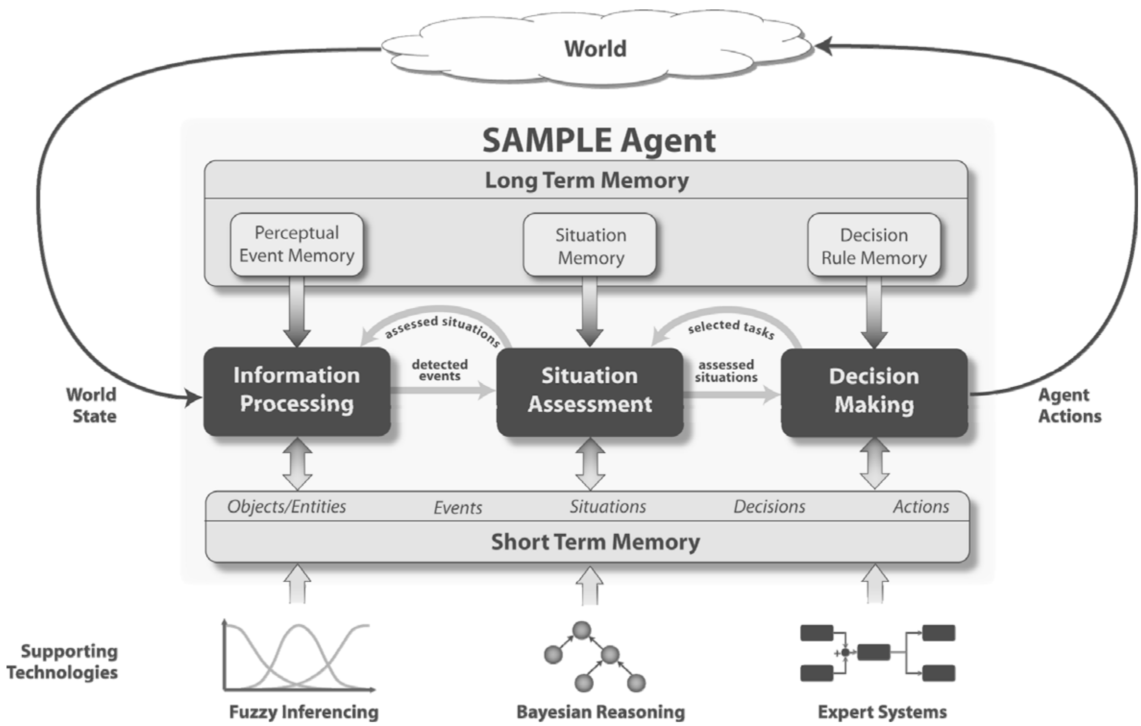

Fig. 1. SAMPLE cognitive modeling architecture.

The processed data, in the form of detected events and structured states, is passed to a *situation assessment* module. This module primarily uses Bayesian reasoning to reason about probabilistic events and situations in an attempt to model the agent's understanding of the unfolding situation through deductive and abductive reasoning. This, in keeping with our recognition-primed approach to behavior modeling, is where the bulk of the complex reasoning and inference are done. We have found that, although Bayesian reasoning is not fully consistent with cognitive experiment data [15], it does provide a good approximation of the human ability to fuse both causal and diagnostic information and meta-information. Other approaches to modeling human situation awareness (e.g., rule-based systems) generally have difficulty modeling such information processing [62].

The assessed situation is passed to a *decision-making* module that uses it to select appropriate responses and actions. SAMPLE uses traditional rule-based expert-system technology [63] for this purpose. In keeping with the recognition-primed decision-making philosophy of SAMPLE [60], the decision-making module in models of expert users is typically relatively straightforward; once the situation is properly understood, experts tend to be able to act fairly directly without the need for more complex planning and problem-solving processes. So, in many situations, this stage can be as simple as a response lookup. In other cases, however, this step performs more complex reasoning in an attempt to satisfy conflicting goals or to respond most appropriately to a complex, uncertain, or unexpected situation.

Each of the three main processing modules has access to both long-term and short-term memory. Long-term memory is used to store the expert knowledge of the agent, such as specific BBNs used for analyzing the domain. Short-term memory is used to store information about the current state of the agent and the environment, such as the particular belief values associated with nodes in the BBN.

3. Bayesian approaches to modeling reasoning about and with meta-information

An important and difficult aspect of modeling human cognitive and behavioral processes is the need to reflect the known impacts of meta-information on those processes. Our analyses across decision-making domains, has revealed a number of different types of qualifiers that can influence information interpretation, understanding, and resulting action, and has made clear the cognitive complexity that should be captured by efforts in any representation of human thought and behavior. HBMs must necessarily process incoming information according to its meta-information, assess the situation represented by that information given its contextualizing meta-information, and select behavioral options based on that assessment. Our prior research on the analysis of the role of meta-information in human reasoning [3,8] has led us to identify five features of this reasoning that need representation within human behavior models:

(1) The model should succeed or fail to *recognize relevant meta-information* as a function of attentional and cognitive demands (e.g., when too much information is present, the model may fail to appreciate that a message is stale and therefore no longer true).
(2) The model should support the representation of successful and unsuccessful human strategies to *process information according to meta-information* (e.g., the model should respond to email tagged as "critical" first).
(3) The model should represent the *aggregation of meta-information* (e.g., the model should be able to fuse meta-information about different sources, different levels of credibility, different time stamps, and different perceived relevance to current activities).
(4) The model should capture *how effectively meta-information is understood* relative to any prior understanding or knowledge (e.g., the model may assume that the situation is dangerous because meta-information about the prior reliability of a warning's source is not factored into the assessment).
(5) The model should succeed and fail at *incorporating meta-information-mediated situation assessments into behavior* or decisions (e.g., the model may understand that the warning's source is not reliable, but still react to a threat).

Our goal in this work is to illustrate methods by which these aspects of human reasoning and decision-making processes with and about meta-information can be represented.

Below, we describe our general approach to modeling using Bayesian belief networks (BBNs) and why this computational formalism is well suited to modeling human reasoning about meta-information (Section 3.1). Next, we discuss how BBNs could be applied to modeling the *recognition*,*processing*, and *aggregation* of meta-information (Section 3.2). Then, we discuss the application of meta-information within models to influence situation assessment (e.g., *understanding*) and decision-making (e.g., *behavior*) (Section 3.3).

3.1. Methods for representing human reasoning

It has been our experience that Bayesian belief networks (BBNs) are particularly versatile tools for modeling a wide range of HBM meta-information reasoning requirements, including the recognition, processing, aggregation, understanding and application of meta-information. In this review, we analyze a range of Bayesian modeling approaches that we have taken to integrate meta-information in previous efforts. We recommend additional research focusing on the application of other technologies to meta-information modeling, including, but not limited to, fuzzy set theory, rule-based production systems, and case-based reasoning. Such research efforts would compliment the work described here, and expand the horizon of human behavior modeling.

Another reason BBNs are particularly applicable in this domain is their versatility in addressing multiple types of modeling requirements. In most problem domains, there are two distinct types of meta-information reasoning that can be captured: deductive reasoning, in which we reason about factors that predict an outcome, and abductive reasoning, in which we reason 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, such as personal expectations and information context, as shown in Fig. 2a) or modeling of the sensor's error likelihood (e.g., reasoning abductively about overall sensor error from various factors affected by that error, such as standard sensor errors (e.g., precision errors) and environmental errors, as shown in Fig. 2b). This allows models to capture the range of ways in which people think about meta-information in their decision-making process.

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 meta-information analysis are explored.

3.2. Modeling the recognition and aggregation of meta-information

Decision-makers recognize, process, and aggregate meta-information in a number of ways. In many cases, a human decision-maker will have to compute meta-information from multiple factors, often including data, meta-data, information, and other meta-information (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. Effective cognitive models must behave similarly.

In modeling human reasoning, data streams are commonly not tagged as information or meta-information. Data and meta-data can map to components of meta-information in a number of different ways, including:

- One-to-one mappings, in which specific data and meta-data components can be directly read as meta-information (e.g., latency of a sensor report).
- Many-to-one mappings, in which a number of data and meta-data components from a number of sensors can be fused into one component of meta-information (e.g., fuse the differences between reports from five sensor streams into one *uncertainty* value).
- One-to-many mappings, in which a number of meta-information components can be extracted from a simple data element (e.g., given a particular sensor type, one can infer reliability and pertinence of the data stream).
- Many-to-many mappings, in which a number of data streams can be fused to produce a number of meta-information components (e.g., eight sensors with varying characteristics have their streams fused by two different algorithms producing recency and confidence values).

Additionally, once meta-information is calculated, it can influence the information gathering, situation assessment, and decision-making process in each of the above manners. Because there is such a wide range of interactions in the calculation and application of meta-information, there are a wide range of uncertainty-oriented technologies that can be used to model it, as discussed in Section 2.2.

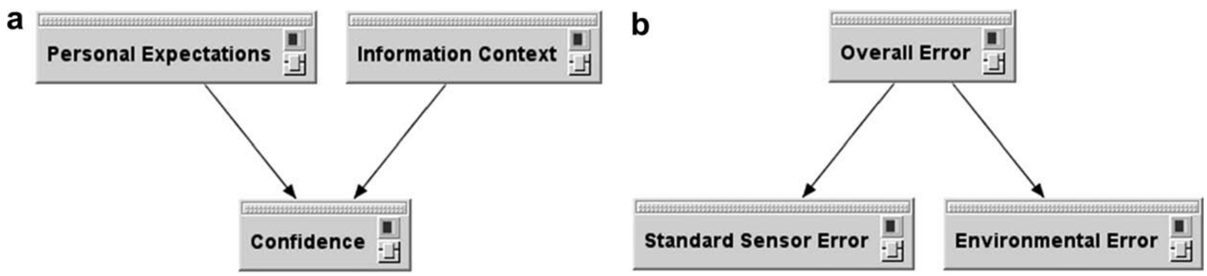

Fig. 2. Deductive vs. abductive reasoning: (a) deductive and (b) abductive.

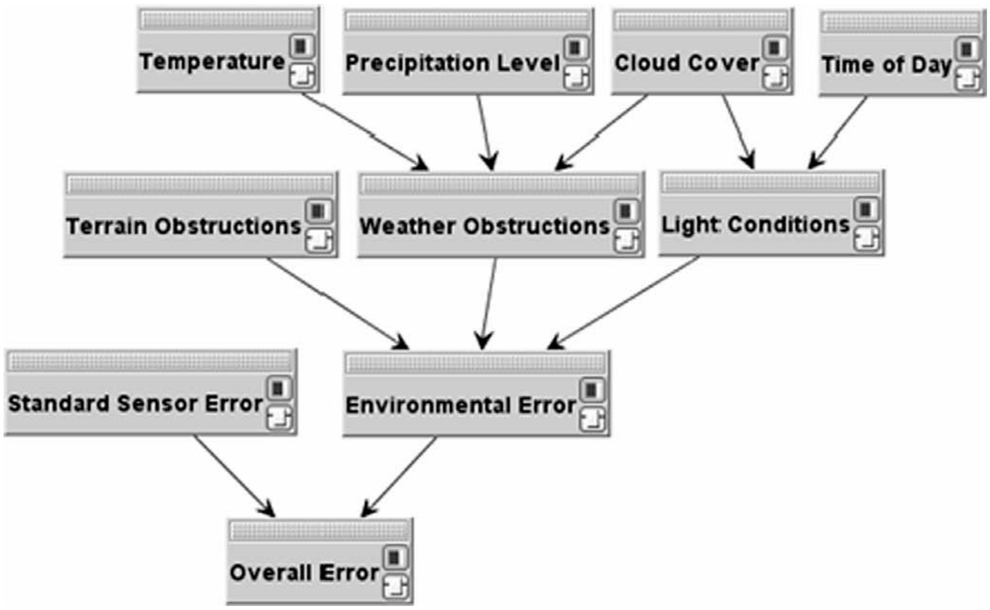

Fig. 3. Computing environmental error meta-information from environmental data.

BBNs, in particular, can be directly applied to model this cognitive computation of meta-information, as illustrated in Fig. 3, which shows a BBN designed to deductively combine various environmental data components with intrinsic sensor error to calculate an overall error for a sensor. In this example, features that lead to certain error conditions are specified, such as *cloud cover*, *temperature*, and *precipitation level* providing evidence for *weather obstructions*, which in turn provides evidence for *environmental errors*.

This BBN does not explicitly include the effect of the sensor type, or other moderating information, on the ways in which these environmental factors influence errors. Additional moderators can be added in several different ways. First, separate BBNs can be used to model different systems (e.g., a second sensor might be unaffected by, and therefore drop, the light conditions branch of this BBN). While this provides the most flexibility in addressing highly complex differences in the models, it can also pose a significant computational detriment, as additional BBNs are loaded and used. Another possibility would be to write software that dynamically modifies the conditional probability tables (CPTs) within the network based on the currently selected sensor (e.g., making one sensor more sensitive to weather changes than another), or to add a node to the network that changes the CPT behavior based on the selected sensor (as shown in Fig. 4).

This reduces the computational complexity, but can reduce the transparency of the BBN's behavior. In more simplistic cases, designers can go so far as to add nodes to the network to manage the new component (e.g., add a node that specifies which sensor type is being used). While this can make a calculation more accurate, it can rapidly increase the complexity of the BBN. Another issue that often arises is the calculation of this meta-information 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). A DBN is an extension of a static BBN used to model a stochastic process (i.e., to model a world that changes over time), by using past and current obtained evidence to compute beliefs about the past, current, and future state of the world [37].

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), different meta-information components will often need to be aggregated. One approach to combining meta-information of various types is aggregation through the application of BBNs. Fig. 5 provides an example of aggregating different types of meta-information, where the effect of available information on confidence (e.g., is there lots of supporting or conflicting information from other data sources?), the effect of current expectations on confidence (e.g., does the given information fit what the human 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 of a specific sensor report, impacting the interpretation of that report.

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 amongst subject matter experts may provide some justification, but will often unavoidably result in inclusion of biases (e.g., predisposition to particular source types or biased interpretation of meta-information influences). On the other hand, using engineering data about sources may not adequately represent how a human would reason about meta-information, resulting in less reflective human behavior models. Therefore, the development of methods that model how decision-makers aggregate meta-information is an open research chal-

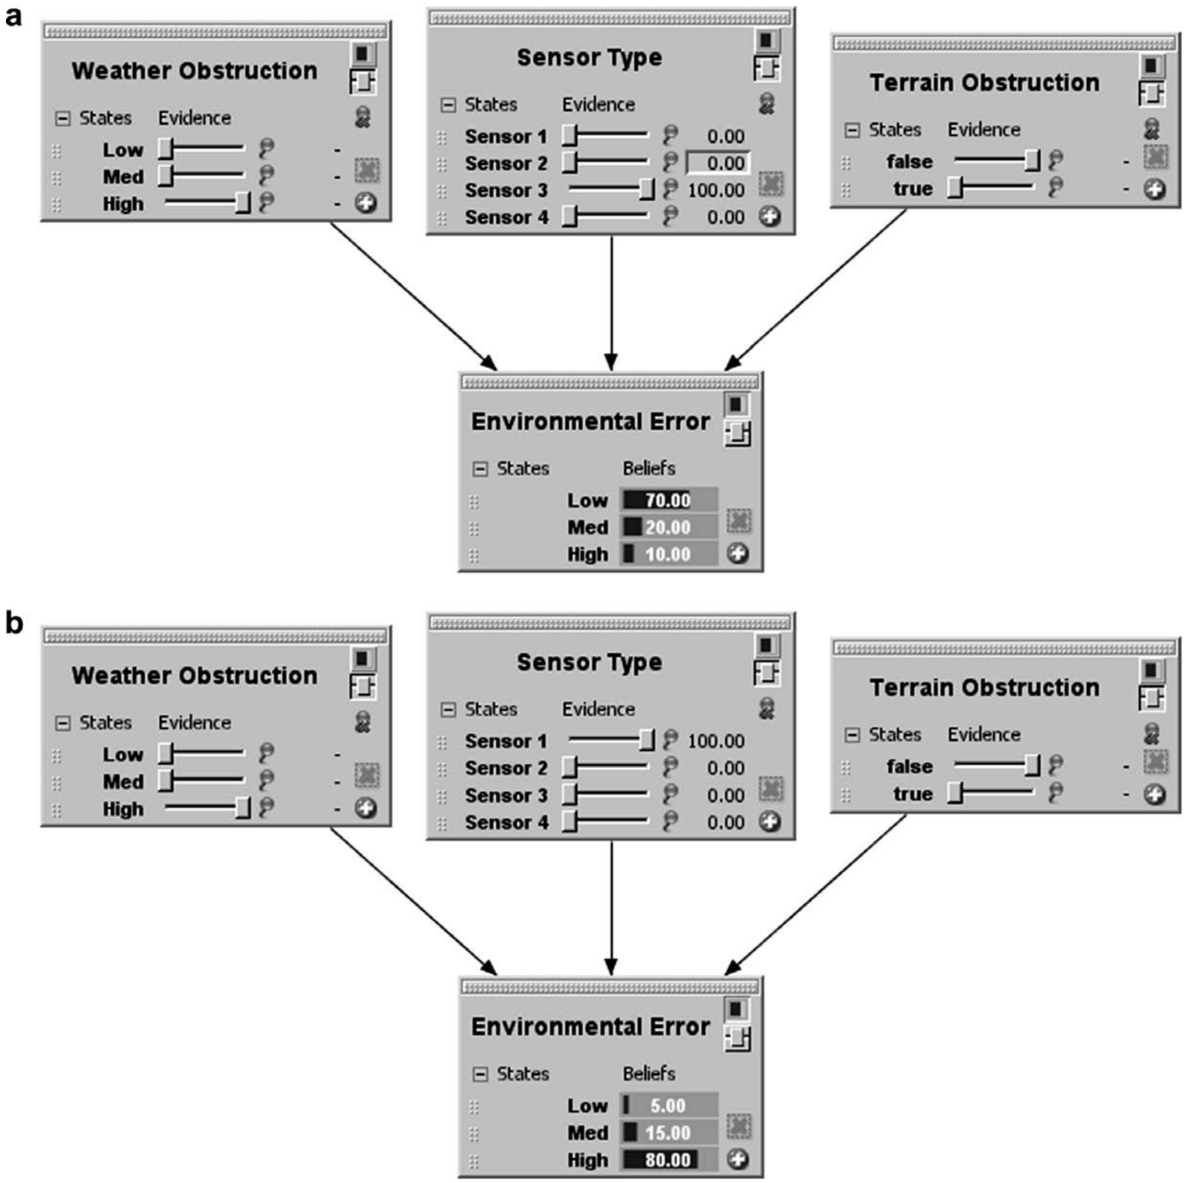

Fig. 4. Sensor type as node in network: (a) sensor 3 more sensitive to terrain obstruction and (b) sensor 1 more sensitive to weather obstruction.

lenge 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.

3.3. Modeling the impact of meta-information on situation assessment

If we assume that meta-information can be computed and/or aggregated in some tractable manner by the cognitive model, then the next step is to capture how meta-information could be used in modeling a human reasoning process, including situation *understanding* and *behavior*. One approach is to simply 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 (e.g., only selecting those with the highest confidence). This approach requires some degree of cognitive task analysis (and/or human-in-the-loop experimentation) to determine how an 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 allocation and filtering mechanisms have already been modeled

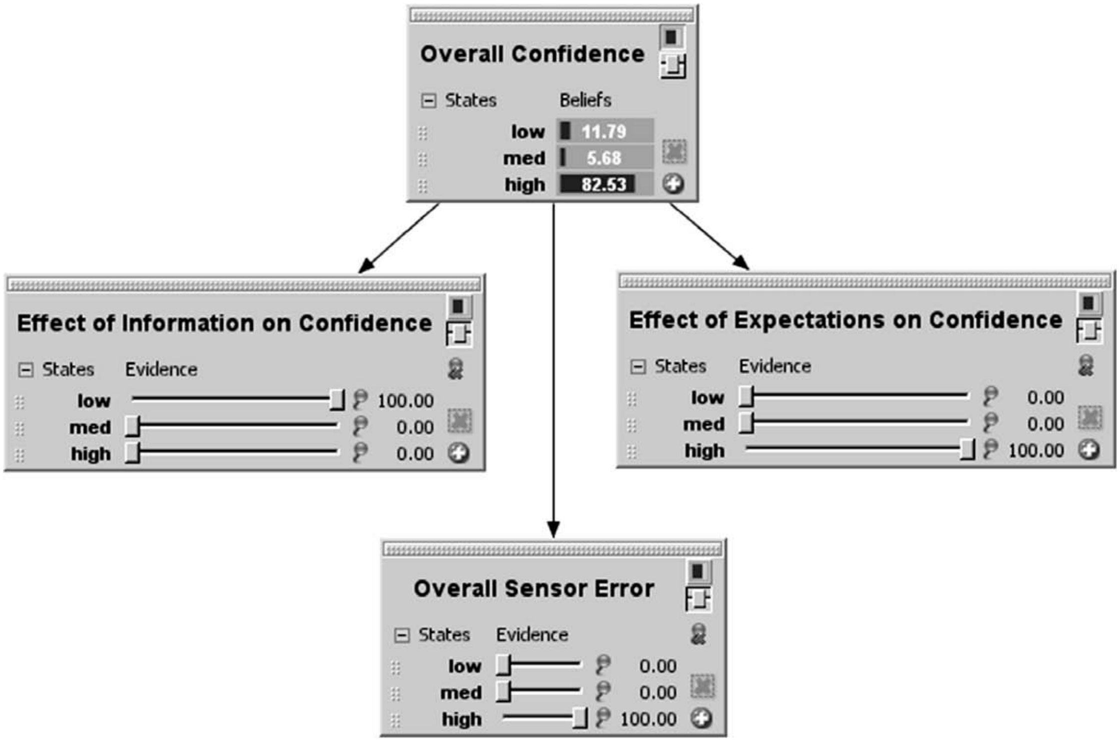

Fig. 5. Aggregating meta-information to compute overall confidence.

in some detail [64], 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 BBN models of information gathering, situation assessment, and decision-making processes. This involves generating BBNs to base beliefs and actions on the application of meta-information (e.g., if a 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 Fig. 6. In this example, when information is posted informing the model that the confidence is low, the threat level increases, representing the expert pilot 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 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 evidence posted on a particular node based on meta-information. We find that by including meta-information in these "glue-code" approaches, we can ensure meta-information has precise effects on data, without dramatically modifying the complexity of our models. Examples of how evidence could be alternatively computed are shown in Table 2. Here, each of the calculated probabilities represents one possible discrete value for a node in a BBN. Rather than setting information values in behavior modeling BBNs based purely on incoming sensor information, the evidence is moderated externally based on meta-information such as sensor reliability, information confidence, credibility, and sensor type, following guidelines extracted from cognitive task analysis. For example, in Table 2, the probability that the location of a detected entity is "near" is calculated by multiplying the sensor value (K) by a meta-information representation of the *reliability* of that sensor, rather than simply processing the sensor value itself. 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.

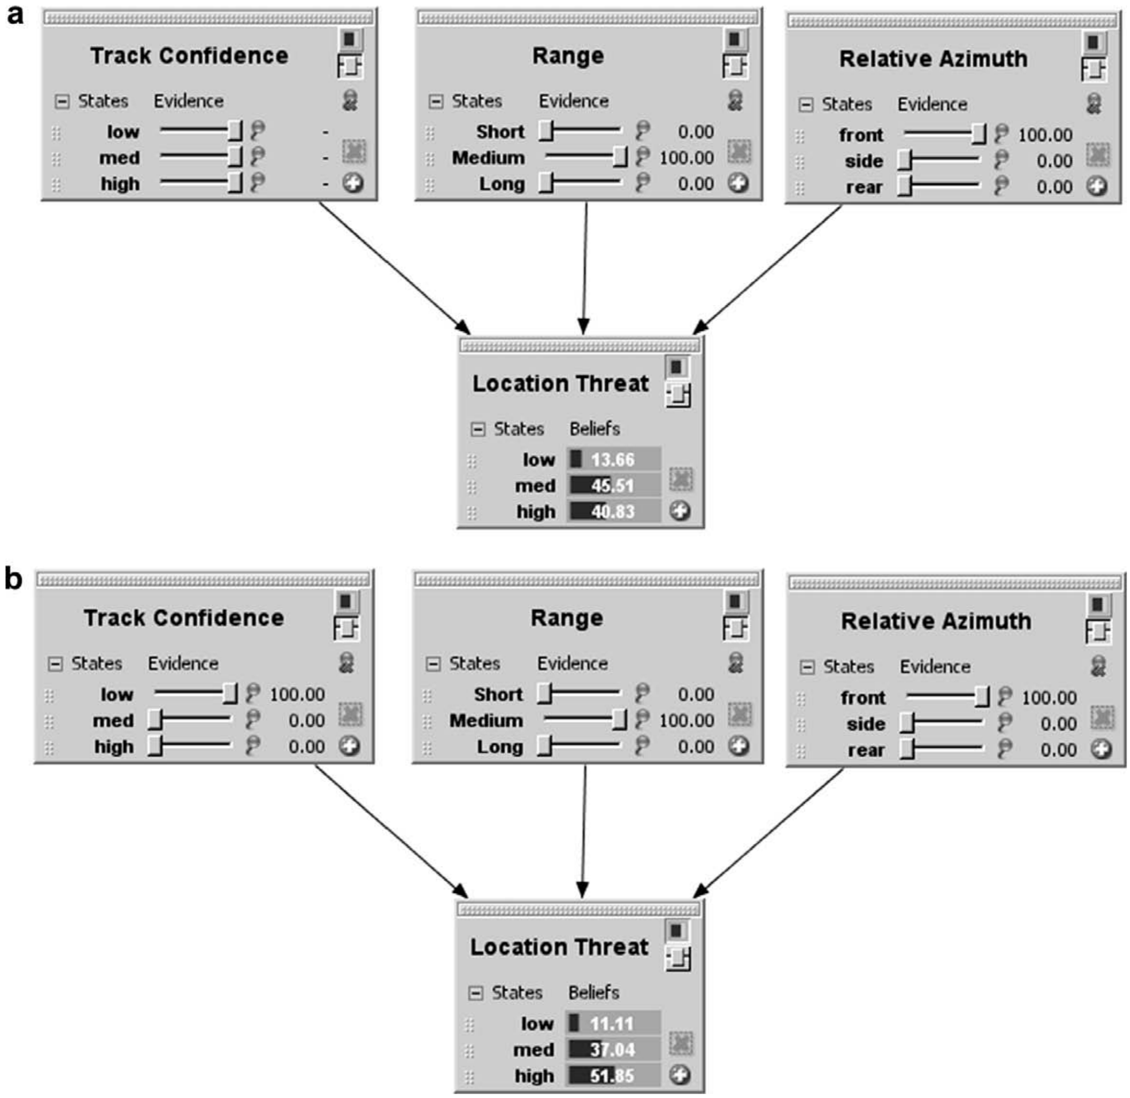

Fig. 6. Incorporating meta-information explicitly into a BBN: (a) no confidence information posted and (b) low confidence posting increases analyzed threat.

Table 2
Examples of computing the probability of a discrete value for a BBN node

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. Conclusion

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 analyses of human cognition and action in different decision-making domains has shown that humans use (and/or fail to 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 and rules 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 meta-information and BBNs in modeling each of the following types of cognitive tasks:

- Recognition of relevant meta-information based on aggregation of available data, meta-data, information, and meta-information into types of meta-information.
- Filtering and prioritization of information based on meta-information.
- Aggregation of different types of meta-information to acquire their combined impact.
- Understanding of the impact of meta-information on existing knowledge.
- Incorporation of meta-information into mediation of situation assessment and decision-making.

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 more general need to more carefully include the influence of meta-information when designing complex human behavior models (and/or systems that support human reasoning [3]). 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 in other modeling architectures.

Acknowledgements

The work described in this paper was performed, in some part, under OSD Contract No. A8650-04-M-6418, US 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

[1] M. Endsley, D. Garland, Situation Awareness, Analysis, and Measurement, Lawrence Erlbaum Associates, Mahwah, NJ, 2000.
[2] M. Endsley, Measurement of situation awareness in dynamic systems, Human Factors 37 (1) (1995) 65–84.
[3] J. Pfautz, E. Roth, A. Bisantz, J. Llinas, A. Fouse, The role of meta-information in C2 decision-support systems, in: Command and Control Research and Technology Symposium, San Diego, CA, 2006.
[4] A. Bisantz, S. Schinzing, J. Munch, The effects of format and level of specificity for displaying uncertainty, Human Factors (2004).
[5] L. Cosmides, J. Tooby, Are humans good intuitive statisticians after all? Rethinking some conclusions from the literature on judgement under uncertainty, Cognition 58 (1996) 1–73.
[6] B. Kuipers, A. Moskowitz, J. Kassirer, Critical decisions under uncertainty: representation and structure, Cognitive Science 12 (1988) 177–210.
[7] J. Pfautz, A. Fouse, M. Farry, A. Bisantz, E. Roth, Representing meta-information to support C2 decision making, in: International Command and Control Research and Technology Symposium (ICCRTS), 2007.
[8] J. Pfautz, E. Roth, A. Bisantz, A. Fouse, S. Madden, T. Fichtl, The impact of meta-information on decision-making in intelligence operations, in: Human Factors and Ergonomics Society Annual Meeting, Orlando, FL, 2005.
[9] E. Roth, E. Patterson, 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, Lawrence Erlbaum Associates, Mahweh, NJ, 2005.
[10] S. Potter, J. Gualtieri, W. Elm, Case studies: applied cognitive work analysis in the design of innovative decision support, in: E. Hollnagel (Ed.), Handbook for Cognitive Task Design, Lawrence Erlbaum Associates, London, 2003.
[11] S. Potter, E. Roth, D. Woods, W. Elm, Bootstrapping multiple converging cognitive task analysis techniques for system design, in: S.F. Chipman, V.L. Shalin, J.M. Schraagen (Eds.), Cognitive Task Analysis, Lawrence Erlbaum, NJ, 2000, pp. 317–340.
[12] K. Harper, S. Mulgund, G. Zacharias, T. Menke, SAMPLE: situation awareness model for pilot-in-the-loop evaluation, in: 9th Conference on Computer Generated Forces and Behavioral Representation, Orlando, FL, 2000.
[13] J. Laird, Soar: an architecture for general intelligence, Artificial Intelligence 33 (1) (1987) 1–64.
[14] J. Anderson, C. Lebiere, Atomic Components of Thought Hillsdale, Erlbaum Associates, NJ, 1998.
[15] C. Wickens, J. Hollands, Engineering Psychology and Human Performance, third ed., Prentice-Hall, Upper Saddle River, NJ, 2000.
[16] C. Lebiere, J. Anderson, D. Bothell, Integrating working memory and perception in a dynamic environment using the ACT-R architecture, in: Symposium on Human Perf. Modeling Evaluation at the XIVth Triennial Congress of the International Ergonomics Association and the 44th Annual Meeting of the Human Factors and Ergonomics Society (HFES/IEA), San Diego, CA, 2000.
[17] D. Meyer, Epic computational models of psychological refractory-period effects in human multi-task performance, vol. TR-94/ONR-EPIC-2, University of Michigan, Ann Arbor, 1994.
[18] R. Marinier, J. Laird, Toward a comprehensive computational model of emotions and feelings, International Conference on Cognitive Modeling (2004).
[19] C. Wallenstein, S. Neal Reilly, K. Harper, B. McNally, Building a human behavior model for a collaborative air-combat domain, BRIMS (2006).
[20] S. Mulgund, G. Zacharias, Situation assessment modeling and pilot state estimation for tactical cockpit interfaces, in: HCI International, San Francisco, CA, 1997.
[21] D. Kieras, J. Rosbe, R. Chong, Modeling with perceptual and memory constraints: an EPIC-Soar model of a simplified enroute air traffic control task, vol. A009614, Soar Technology Inc., Ann Arbor, MI, 2001.
[22] K. Harper, S. Ho, G. Zacharias, Intelligent hostile urban threat (IHUT) agent for MOUT operations, in: 9th Conference on Computer Generated Forces and Behavior Representation, Orlando, FL, 16 May 2000.
[23] K. Harper, S. Guarino, A. White, M. Hanson, K. Bilimoria, D. Mulfinger, An agent-based approach to aircraft conflict resolution with constraints. AIAA Guidance, Navigation, and Control, Monterey, CA, 2002.

[24] P. Aykroyd, J. Bachman, K. Harper, E. Hudlicka, Modeling individual differences and stressors using the SAMPLE cognitive architecture, in: 14th Conference on Behavior Representations in Modeling and Simulation, Universal City, CA, 2 May 2005.
[25] S. Trickett, G. Trafton, L. Saner, C. Schunn, I don't know what is going on there: the use of spatial transformation to deal with and resolve unc. in complex vis, in: M. Lovett, P. Shah (Eds.), Thinking with Data, Erlbaum, Mahwah, NJ, 2005.
[26] N. Sarter, B. Schroeder, Supporting decision making and action selection under time pressure and uncertainty, Human Factors 43 (2001) 573–583.
[27] E. Watkins, Improving the Analyst and Decision-Maker's Perspective through Uncertainty Visualization, Master of Science, Air Force Institute of Technology, Department of Systems and Engineering Management, 2000.
[28] A. Bisantz, Modeling Environmental Uncertainty to Understand and Support Dynamic Decision Making, School of Industrial and Systems Engineering, Georgia, IT, 1997.
[29] K. Teigen, The language of uncertainty, Acta Psychologica 68 (1988) 27–38.
[30] D. Kahneman, P. Slovic, A. Tversky, Judgment under Uncertainty: Heuristics and Biases, Cambridge University Press, Cambridge, UK, 1982.
[31] C. Schunn, S. Kirschenbaum, J. Trafton, The ecology of uncertainty: sources, indicators, and strategies for information uncertainty, 2003. <http://www.au.af.mil/au/awc/awcgate/navy/nrl_uncertainty_taxonomy.pdf>.
[32] M. Yovits, J. Abilock, A semiotic framework for information science leading to the development of a quantitative measure of information, in: 37th American Society for Information Sciences Meeting, vol. 11, 1974, pp. 163–168.
[33] R. Lipshitz, O. Strauss, How decision-makers cope with uncertainty, in: Proceedings of the Human Factors and Ergonomics Society 40th Annual Meeting – 1996, Human Factors Society, Philadelphia, 1996, pp. 189–193.
[34] G. Klein, Sources of Power: How People Make Decisions, MIT Press, Cambridge, MA, 1998.
[35] J. Halpern, Reasoning about Uncertainty, MIT Press, Cambridge, MA, 2003.
[36] S. Parsons, Qual. Meth. for Reas. under Unc, MIT Press, Cambridge, MA, 2001.
[37] P. Russell, P. Norvig, Artificial Intelligence: A Modern Approach, second ed., Prentice-Hall, Upper Saddle River, NJ, 2003.
[38] J. McCarthy, P. Hayes, Some philosophical problems from the standpoint of artificial intelligence, in: B. Meltzer, D. Mitchie, M. Swann (Eds.), Machine Intelligence, vol. 4, Edinburgh University Press, Edinburgh, Scotland, 1969, pp. 463–502.
[39] F. Lin, An argument-based approach to non-monotonic reasoning, Computational Intelligence 9 (1993) 254–267.
[40] L. Zadeh, Fuzzy sets, Information and Control 8 (1965) 338–353.
[41] J. Pearl, Causality: Models, Reasoning, and Inference, Cambridge University Press, 2001.
[42] D. Geiger, D. Heckerman, Knowledge representation and inference in similarity networks and Bayesian multinets, Artificial Intelligence 82 (45) (1996) 74.
[43] R. Howard, J. Matheson, Influence diagrams, in: R. Howard, J. Matheson (Eds.), Readings on the Principles and Applications of Decision Analysis, Strategic Decisions Group, Menlo Park, CA, 1984, pp. 719–762.
[44] M. Borsotto, W. Zhang, E. Kapanci, A. Pfeffer, C. Crick, A junction tree propagation algorithm for Bayesian networks with second-order uncertainties, in: 18th IEEE Conference on Tools with Artificial Intelligence (ICTAI'06), 2006, pp. 455–464.
[45] P. Walley, Stat. Reasoning Imprecise Prob., Chapman & Hall, London, NY, 1991.
[46] P. Smets, Imperfect information: imprecision and uncertainty, in: A. Motro, P. Smets (Eds.), Uncertainty in Information Systems, Kluwer, Boston, MA, 1997, pp. 225–254.
[47] M. Smithson, Ignorance and Uncertainty: Emerging Paradigms, Springer Verlag, New York, NY, 1989.
[48] P. Bosc, H. Prade, An introduction to the fuzzy set and possibility theory-based treatment of flexible queries and uncertain or imprecise databases, in: A. Motro, P. Smets (Eds.), Uncertainty in Information Systems: From Needs to Solutions, Kluwer, Boston, MA, 1997, pp. 285–324.
[49] C. Elkan, The paradoxical controversy over fuzzy logic, IEEE Expert 9 (4) (1994) 47–49.
[50] H. Havenstein, Metadata management returns to the fore, Computer World, 18 July 2006.
[51] D. Marco, M. Jennings, Universal Meta-Data Models, Wiley, New York, NY, 2004.
[52] J.M.C. Schraagen, S.E. Chipman, V. Shute, J. Annett, M. Strub, C. Sheppard, J.Y. Ruisseau, N. Graff, State-of-the-Art Review of Cognitive Task Analysis Techniques, vol. TNO Report TM-97 B012, TNO Human Factors Research Institute, Soesterberg, The Netherlands, 1997.
[53] C. Burns, J. Hajdukiewicz, Ecological Interface Design, CRC Press, 2004.
[54] W. Elm, S. Potter, J. Gualtieri, E. Roth, J. Easter, Applied cognitive work analysis: a pragmatic methodology for designing revolutionary cognitive affordances, in: E. Hollnagel (Ed.), Handbook for Cognitive Task Design, Lawrence Erlbaum Associates, London, 2003.
[55] M. Endsley, B. Bolte, D. Jones, Designing for Situation Awareness: An Approach to Human-Centered Design, Taylor & Francis, London, 2003.
[56] K. Vicente, Cognitive Work Analysis: Towards Safe, Productive, and Healthy Computer-based Work, Lawrence Erlbaum Associates, Mahwah, NJ, 1999.
[57] J. Rasmussen, A. Pejtersen, L. Goodstein, Cognitive Systems Engineering, Wiley and Sons, New York, 1994.
[58] R. Pew, A. Mavor, Modeling Human and Organizational Behavior: Application to Military Simulations, 1998.
[59] G. Klein, Recognition-primed decisions, Advances in Man-Machine Systems Research 5 (1989) 47–92.
[60] G. Klein, Naturalistic decision making: implications for design, Gateway 5 (1) (1994) 6–7.
[61] L. Zadeh, Outline of a new approach to the analysis of complex systems and decision processes, IEEE Transactions on Systems, Man and Cybernetics 3 (1) (1973) 28–44.
[62] J. Pearl, Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference, Morgan Kaufmann, San Mateo, CA, 1988.
[63] E. Shortliffe, Computer-Based Medical Consultations: MYCIN, American Elsevier, New York, NY, 1976.
[64] K. Harper, G. Zacharias, Modeling attention allocation and multitasking in computational human behavior representations, in: 11th Conference on Computer Generated Forces and Behavior Representation, Orlando, FL, 7 May 2002.