12th ICCRTS Abstract Submission
"Adapting C2 to the 21st Century"

Representing Meta-Information to Support C2 Decision Making

Cognitive and Social Issues, Modeling and Simulation, C2 Technologies and Systems

DRAFT PAPER
Submission Number: I-188

Jonathan Pfautz
Adam Fouse
Michael Farry
Ann Bisantz
Emilie Roth

Corresponding Author:
Jonathan Pfautz
Charles River Analytics
625 Mt. Auburn Street, Cambridge, MA 02138
(617) 491-3474 / (617) 868-0780
jpfautz@cra.com

Abstract:

Aggregating, assimilating, and understanding the ever-larger amounts of heterogeneous information present in network-centric environments presents distinct cognitive challenges to the command and control staff. Under previous efforts (Pfautz, 2006), we have detailed our efforts to analyze how qualifiers of information, or *meta-information* (e.g., uncertainty, recency, pedigree), impact information processing and situational awareness in an already challenging decision-making environment. To date, few existing systems explicitly support the management and representation of meta-information. Here, we describe several specific efforts to develop methods for the representation of meta-information in C2 decision-support tools, including methods to support asset allocation (e.g., for air-based ISR, for addressing ground-based threats, for neutralizing near-space or space-based threats). These methods include techniques for the visual portrayal of meta-information in C2 decision-making systems as well as approaches to the computation, when necessary, of that meta-information. In this paper, we discuss these methods within example domains, and discuss lessons learned for the design of future C2 decision-support systems.

1. Introduction

Modern military operations rely on the ability of command and control personnel to manage large volumes of dynamic information from heterogeneous sources. The metamorphosis of a traditional military information environment into a network-centric environment has increased this management burden as more and more data needs to be processed, aggregated, and understood to provide actionable intelligence. This burden is exacerbated by advances in sensor technology that result in an increase in the types and sources of data, and advances in information technology that results in an increase the types and degrees of processing of that data. The combination of these challenges results in known impacts on the human trying to process and understand information (i.e., their "situational awareness" (Endsley & Garland, 2000; Endsley, 1995)), such as increased workload (Endsley, 1993) information overload (Woods, Patterson, & Roth, 2002), and attentional blindness (Durlach, 2004).

Further complicating C2 decision-making in these overwhelming information environments, is the need to perceive, understand, and reason about how information from heterogeneous sources (from both sensors and sensor fusion systems) is qualified. The qualifications of information, or its *meta-information* (e.g., pedigree, recency, uncertainty), provides context to the information in a specific decision-making task; therefore the presence or absence of meta-information can critically influence how a decision-maker will process, understand, and act on that information (Pfautz et al., 2005c). Our analysis of cognitive tasks across different domains (e.g., military command and control, intelligence analysis, sensor management, weather impact analysis, wildfire management, among others) has revealed the clear role of meta-information in decision-making (Pfautz et al., 2006b). As such, we have developed the following working definitions for terms that we use throughout this paper (Pfautz et al., 2005b; Potter et al., 2000):

- *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, meaningful, and/or useful in a directed decision-making process or resulting 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 +/- 1.5m)
- *Meta-information* is characteristics or qualifiers of *information*, affecting a human's:
  - 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 hours old, the model must actively gather new information before moving into that location)

As noted in our prior work, categorizing inputs according to these definitions (i.e., mapping data and meta-data to information and meta-information) is dependent on the particular cognitive task and the context in which that task is performed. Nevertheless, these definitions serve to explicitly identify the critical role of meta-information in human decision-making.

Understanding of the role of meta-information in a particular C2 domain and within a specific cognitive task is critical to developing systems to support C2.

Such an understanding, and its impacts on the design and development of support systems is facilitated by the application of *Cognitive Systems Engineering* (CSE), a methodology for defining aspects of human reasoning and behavior to aid system design (Schraagen, Chipman, & Shalin, 2000; Rasmussen, Pejtersen, & Goodstein, 1994). CSE involves iterative several phases of analysis, prototype development, and design. The analysis phase can be accomplished with a variety of methods, including: Cognitive Task Analysis (Potter et al., 2000), Cognitive Work Analysis (Vicente, 1999), Goal-Directed Task Analysis (Endsley & Rodgers, 1994), or Applied Cognitive Work Analysis (Elm et al., 2003). Each of these analysis methods attempts to uncover, through formal observation and interviewing techniques, the salient characteristics of the work domain as well as the cognitive processes used to perform work. In analysis efforts described in previous work (Pfautz et al., 2005b), we identified the main types of meta-information that impact the decision-making process within a set of well defined application domains.

Across these analysis efforts, we also developed prototype concepts for decision-support, and conducted user 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 (across the different domains) and over 500 hours of interviews, observation, and evaluation, and has resulted in the development of a list of sources and types of meta-information we have consistently encountered across application domains (Table 1).

Table 1: Sources and Types of Meta-Information in Explored Domains (adapted from Pfautz et al., 2006)

Meta-Information Type	Sub-Types or Related Types
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 & history Directly observing or deriving information
Characteristics of the source as a function of other factors	Time Location in environment (e.g., terrain, weather) Types of intermediate processing Content of report Information context
Uncertainty	Spatial uncertainties Temporal uncertainties Uncertainties about uncertainty reporting Likelihood Probability Confidence Accuracy Precision
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 Frequency of reporting of information Missing or degraded information qualifiers Information-to-noise ratio History
Reliability of source	W.r.t. source characteristics W.r.t. information context
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

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
Temporal qualifiers	Staleness Recency Certainty about time of reporting Latency Lag Absence of expected information

While these sources and types of meta-information were necessarily specific to the domains we studied, they provide the basis for exploring new C2 domains and developing new support systems for those domains. In this paper, we will expand upon this theoretical analysis of meta-information to show the application of these concepts to system and display design.

While some of our earlier efforts to develop systems to accommodate meta-information requirements are described elsewhere (Bisantz et al., 2006; Lefevre, Pfautz, & Jones, 2005; Pfautz et al., 2005a; Pfautz et al., 2005b), we present here exemplars of our latest efforts to understand meta-information and its impacts on system design. In our recent work, we have developed prototype systems to support reasoning with meta-information, and as part of these systems have designed specific computational methods to support generation and refinement of meta-information and visualization methods to effectively represent meta-information to a human decision-maker. The examples presented here, drawn from work across several domains, serve to elucidate the concepts described previously as well as identify future research and development needs.

2. Background

Because this paper describes approaches to supporting reasoning about meta-information with both computation and visualization, we discuss below the applicable literature. Unfortunately, most literature addressing computation and portrayal of meta-information are focused on only one type of meta-information, namely *uncertainty*. In our work, described in Section 3, we build upon many of the described techniques and attempt to extend them to address meta-information more generally.

2.1 Computational Methods and Meta-Information

Computational methods are inherent to a huge variety of C2 systems. To date, these systems have tended to focus on methods to reason about uncertainties present in the real work in tasks such as monitoring plan execution and anticipating adversary courses of action. There are a variety of computational approaches that have been developed by the Computer Science, Artificial Intelligence, Signal Processing, and Data Fusion communities to support reasoning about one or more types of uncertainty (for an extensive treatment of the subject, see (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; Pearl & Russell, 2000), 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 or semi-automated reasoning about uncertainty, and a set of techniques that may be of use to provide computational support for reasoning about meta-information.

Another representational mechanism that may prove useful in designing systems for reasoning about and with large volumes of heterogeneous information and associated meta-information is the concept of *ontologies*. Generally, an ontology explicitly defines a conceptual structure in a particular domain (Gruber, 1993). The study of ontologies has moved in recent years from an issue of mainly philosophical concern (Quine, 1969) to an area with wide applications in science and engineering. The use of ontologies to structure data and meta-data allows the application of computational methods designed around those structures, which ideally reflect the critical characteristics of the domain. Knowledge obtained from domain experts, encoded in an ontology, and embodied in a computational method could be used to provide automated reasoning (e.g., experts describe the key characteristics that differentiate benign aircraft from incoming missiles; an ontology is developed to capture the key thresholds for each characteristic; and a computational method is developed to automatically detect, based on the ontology and the expert knowledge, if an incoming aircraft is a threat). The use of ontologies (or other approaches to providing formal data structures across humans and machine) may be of great utility in developing methods for the computation of, or automated reasoning about, meta-information. Meta-information types, obtained through a cognitive task analysis in an application domain, could be encoded into an ontology, which could then be used with associated computational mechanisms to generate, aggregate, and process meta-information.

2.2 Visualization Methods and Meta-Information

Some recent work has begun to address the need for representing meta-information in display systems. This work includes efforts to design visualization methods for the presentation of multiple types of uncertainty and time-related meta-information for weather-impact analysis (Lefevre et al., 2005). In addition, methods for presenting meta-information in geographic information systems (GIS), particularly those relying on modeling physical phenomena, have been hypothesized and prototyped (Pfautz et al., 2005a). Some of these methods have been developed as untested proof-of-concepts (Pfautz et al., 2006a), others have been empirically evaluated to assess how presentations of meta-information are perceived and understood (Bisantz et al., 2006). All of these efforts to design and develop visualization methods for meta-information have been based, to some extent on prior work on the portrayal of uncertainty.

In the scientific visualization and GIS community, researchers have been investigating the presentation of uncertainty. For example, these researchers hypothesized that certain visual elements could be used to code the uncertainty in large data sets – elements such as attributes of scene geometry (e.g., color, shading, and bumpiness (Pang, Wittenbrink, & Lodha, 1997b)) or traditional graphic variables used in cartography (e.g., texture, color, orientation, and shape (MacEachren, 1992)). Others have postulated the use of use of glyphs (graphical forms such as arrows or vertical lines), in part because it frees other graphical dimensions for other purposes. Glyphs have been used to represent magnitude and direction of winds and ocean currents along with the uncertainties in these dimensions. In one example, the general shape of the glyph was an arrow, with the width of the arrowhead representing uncertainty in heading, and multiple arrowheads representing uncertainty in magnitude (Pang, Wittenbrink, & Lodha, 1997a; Lodha et al., 1996). More recent work on glyphs (Whitaker & Thomas-Meyers, 2006), suggests that these forms can be used to manipulation the context in which information is presented (and therefore we posit that these methods may be used for the representation of meta-information).

3. Representation of Meta-Information for Decision Support

In the following paragraphs, we describe recent work done on a number of different efforts to develop decision-support systems for C2. In all of these efforts, we have applied our CSE approach to system development to identify cognitive tasks and the impacts of meta-information, then prototyped methods to support those tasks with specific computational and visualization methods. In the following examples, we present a rather generic C2 scenario based on the allocation of assets as being representative of the domains in which we performed our research. The methods we present below should be considered to be early design concepts that will be refined as part of our ongoing human-in-the-loop evaluation of these concepts.

3.1 Identification of Sources and Types of Meta-Information

The identification of sources of data and meta-data gives form to the input that is available for meta-information computation methods. These sources range from modeling and simulation environments, to raw input from sensor networks, to user-vetted input. An understanding of useful output (i.e., information and meta-information) is also critical to the design of computation methods. Our team primarily uses knowledge elicitation and requirements analysis, as part of our cognitive systems engineering approach, to identify types of information and meta-information. Applying these techniques across a variety of domains has yielded distinct decision tasks engaging the user, and information and meta-information types which affect the user's decision-making process.

To support these techniques, we construct vignettes in each domain that are representative of decisions that users in that domain are likely to face, the factors that are likely to influence their decisions, and the information and meta-information that needs to be presented to support decision-making. Some key classes of decision types included:

- Decisions related to classification & identification of threats
- Decisions related to optimizing unit movement paths over a specified area
- Decisions related to assessment of the health and status of the units
- Decisions related to timing of action
- Decisions related to threat and risk avoidance through appropriate resource allocation

For each decision, we identified information that influences the outcome. This information included factors such as an Intelligence, Surveillance, and Reconnaissance (ISR) unit's classification of a potential threat, the risk associated with inaction given a threat, the relevant geospatial, temporal, and environmental information, and the current schedule of tasks performed by units.

The tactical decisions in the vignettes represented a diverse array of tasks. For example, one task involved deploying UAVs to scout potential river crossings. Another task considered a resupply operation involving reassignment of an already busy unit. A different task involved tracking the movement of a ballistic missile as it sought out its target. Generally, these decisions focused on the analysis of unit suitability with respect to the task, resource utilization, and scheduling. Other decisions require the decision-maker to consider how a group of units may be deployed to achieve several simultaneous goals.

Meta-information can be used to define the level of uncertainty or the bounds of knowledge of a piece of functional information that is needed to make a decision. For each domain and task, we sought to list the elements of meta-information that would need to be represented to facilitate decision-making. As an example, tactical decisions related to resource management for ISR and resupply tasks involve meta-information with respect to unit capabilities, current unit tasking, tasks currently unfulfilled, current ISR coverage, potential and known threats, and environmental factors. This includes the following types of meta-information:

- Certainty of current unit capabilities (e.g., how certain is it that a UAV has the necessary supplies to complete the resupply task?)
- Vulnerability of a unit to a certain type of threat (e.g., certain types of aerial units may be susceptible to a threat that can fire machine gun rounds)
- Certainty of unit capability with respect to environmental conditions (e.g., a tactical UAV may be more severely affected by wind than a high-altitude UAV, but it depends largely on wind direction)
- Priority of current unit tasks and tasks currently unfulfilled (e.g., 5 separate missile warheads are heading towards different targets, but there are not enough resources to fully monitor them all)
- Certainty of ISR coverage (e.g., a UAV detected a vehicle 20 KM north of the forward operating base, but it may have been a civilian vehicle)
- Recency of ISR coverage
- Certainty of threat location and threat type
- Certainty and recency of weather reports
- Certainty about conflicting paths for aerial units

3.2 Computational Methods and Meta-Information

For each separate task within the developed vignettes, the decision-maker only considers a distinct subset of the above meta-information types. By examining the *specific factors* contributing to decisions that arose in the scenario, we were able to abstract a generalized ontology of factors that contribute to the consideration of meta-information or define the bounds of knowledge for different classes of functional information. These constitute sources of meta-information that need to be represented in decision-support systems, both in the underlying formal data representations and in the visual representations.

The use of ontologies allows the employment of a specific set of meta-information management and computation methods in an operating environment. The ontology also catalogs types of information and meta-information that are relevant in a given situation. The structure of each ontology is a direct result of our efforts in cognitive systems engineering. The information and meta-information requirements of the decision-maker, which are generated from the analysis of decision tasks, populate the ontology. Therefore, the ontology is unique and tailored to a specific setting. At the same time, the use of swappable ontologies within a software framework engenders flexibility – ontologies for wildly different domains can be applied as appropriate (e.g., a ballistic missile defense ontology versus a UAV path management ontology).

There are also practical matters affecting meta-information computational methods. Users must have faith that the information and meta-information being synthesized and presented are relevant to the task at hand, and have been calculated in a sensible way. Therefore, transparency is important: a user should have seamless access to the inner workings of the computation methods. Visualizations of computation, as well as the organization of computation methods into ontologies, are appropriate means to achieve transparency and user trust. To be useful in a decision aid, computation methods must also be able to handle relatively large data sets and produce relevant information and meta-information quantities in a short amount of time. Modern computational technology has advanced rapidly, but this progress is a double-edged sword: data can be processed more quickly, but a much larger body of data is available to process. Great care must be taken in the design of the meta-information computation algorithms, to ensure that relevant information is delivered to the user in a timely fashion.

One computational method that has shown some promise for supporting automated reasoning with and about meta-information is Bayesian Belief Networks (BBN) (Guarino et al., 2006). They provide a graphical representation of causal relations that provides a degree of the required transparency, while remaining general enough to be easily applied to different application domains. In Figure 1 below, we show a BBN developed to compute the certainty of the classification and location of a potential threat. The object may have been identified by a number of friendly forces, each reporting a slightly different location. Each report may have an inherent likelihood to be correct based on the friendly force's ISR equipment, and the distance from the friendly force to the sensed entity. In addition, friendly forces may be providing corroborating or conflicting reports about the entity's classification or status. Since so many factors contribute to the certainty about the threat's location and classification, it is useful for the user to have a reliable metric to gauge an overall certainty rating. The use of BBNs enables the user to see exactly how the certainty rating was derived, and how different combinations of posted evidence may affect the outcome. The exact evidence values are based on parametric values (i.e., user-defined thresholds for "recent" and "not recent", "low" and "high" number of sensor reports, etc.) In this way, users who are not pleased with the current calculation can adjust specific thresholds to achieve a more satisfactory result.

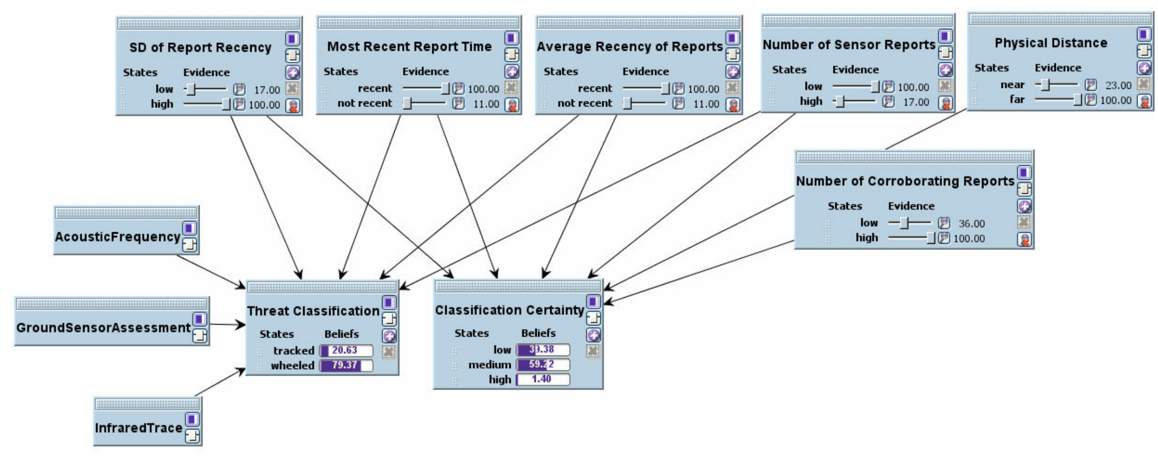

Figure 1: Using a Bayesian Belief Network to calculate certainty of threat calculation

3.3 Supporting Iterative Design and Development

The iterative design and development of visualization systems has been supported by a generic toolkit which provides a framework for the rapid construction of meta-information computation and visualization methods. The ontological structure of the meta-information computation module, and the associated grouping of specific methods, provides an easy affordance for allowing different algorithms to be applied to different domains. The flexibility of the toolkit's visualization creation module allows users to map meta-information quantities to specific visualization techniques, and provides a capacity for the rapid prototyping and evaluation of those techniques.

The computation module is composed of a framework that collects and aggregates reports asynchronously from a variety of data sources, and dynamically provides fresh models of information and meta-information that correspond to specific types of information and meta-information. Within this framework, a collection of individual, interchangeable computation methods are employed in combination, each performing a calculation of limited scope and enhancing the informational model. The appropriate combination is determined by the meta-information ontology selected for the situation. Different user tasks and meta-information requirements call for appropriate suites of computation methods (e.g., a calculation of the certainty of an entity's location, based on data obtained from friendly sensors, may not be useful if said entity is a friendly infantry unit equipped with a fully functional GPS.) The limit on each computation method's responsibility allows different combinations to be evaluated quickly and applied to certain situations, subject to information requirements and the user's requests. These individual methods range from simple mathematical translations (e.g., latitude and longitude coordinates being translated to an appropriate format), to more nuanced methods that extract specific types of meta-information (e.g., calculating the certainty of entity classification through Bayesian reasoning on an unspecified number of sensor report data and the associated meta-data.) Restricting each method's scope also allows for the easy replacement of each method, which in turn facilitates evaluation.

The visualization creation module of the toolkit provides the developer with tools to facilitate the rapid prototyping of meta-information visualization methods. These tools include the ability to control what elements of the incoming data to represent, the ability to map the meta-information to some visual representation method, and the ability to incorporate that representation into an existing display. Various user interfaces have been developed to allow the developer to create templates of mappings between meta-information and visual representations. In this way, the visualization developer can explore combinations of representations when types of meta-information are interdependent (e.g., stale data may indicate one type of uncertainty, while an unreliable source may indicate another type of uncertainty) as well as store, reuse, and edit effective schemas.

3.4 Visualization Methods and Meta-Information

While many existing information visualization techniques are well established (e.g., consistency with natural mappings (Norman, 1988)), few effective meta-information visualization methods have previously been documented (Bisantz et al., 2006; Pfautz et al., 2006b; Pfautz et al., 2005a; Pfautz et al., 2005b). The manipulation of basic visual variables such as transparency or saturation can be employed, seeking to achieve an intuitive understanding of levels of meta-information, such as uncertainty or latency; while the manipulation of hue from green to red can convey a change from positive to negative factors. Other graphical variables may also allow a natural, intuitive mapping (e.g., degree of blurriness for uncertainty of location), while some may require more training (e.g., the use of varying graphical textures, animation, or annotation). Figure 2 shows an example of using transparency to represent meta-information. In this case, a colored area around a UAV on a map represents the coverage of the UAV's infrared sensor. When the sensor is focused on a smaller area, it has low latency, and so the sensed area is represented as less transparency. When the sensor's focus is expanded, its latency increases, and the visual representation increases in transparency.

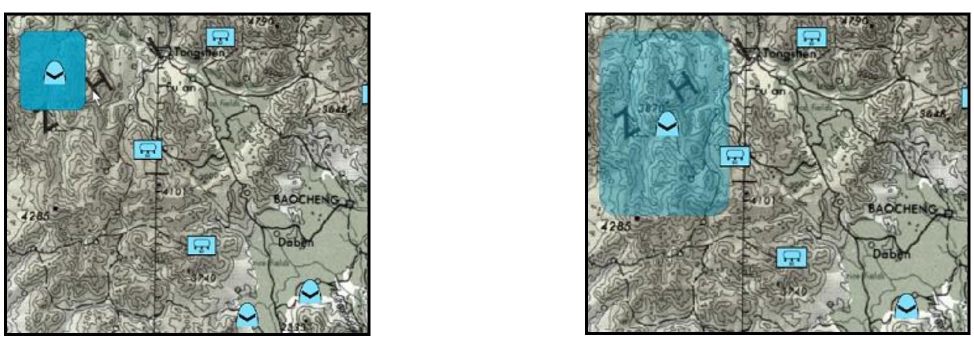

(a) (b)

Figure 2: Using transparency to represent (a) low latency of sensor coverage and (b) high latency of sensor coverage

We have also developed preliminary representations of meta-information in the form of glyphs. These representations were informed in part by research on "knowledge glyphs" that has shown that visual iconic representations can present rich information but have limited dimensionality (Whitaker et al., 2006). To best utilize glyph representations, we have developed interfaces to allow the decision-maker to interact with the level of detail that is presented visually. Figure 3 shows an example of this interface. In this example, the interface has been designed to support tasking of UAVs, and a particular UAV is shown as being inappropriate for the task, but with low certainty, through the use of a transparent highlight (Figure 3a). The decision-maker can increase the level of detail of the representation, revealing that the low certainty derives from information about the weapons system (Figure 3b).

(a) **(b)**

Figure 3: Drill-down glyph representations: (a) user interface for level of detail control and (b) further detail about appropriate asset selection

We have also investigated how to best present meta-information and manipulate its discriminability, visibility, conspicuity, and/or salience as a function of the perceptual and cognitive context. We have exploited knowledge of human attentional mechanisms and visual search (Wolfe, 1998) to determine the parameters by which we can improve the efficiency with which information that is more "important" with respect to its meta-information can be perceived by the user. This approach requires an understanding of human attentional capacity and its impact on visual search (Huang & Pashler, 2005) as studied with various eye tracking methods (Fleetwood & Byrne, 2003). In particular, we have developed methods to manipulate the color difference between an entity and its background, using the color difference number (Wyszecki & Stiles, 1982) as a rough estimation for salience. We are currently investigating the effects of salience manipulation through color adjustment on meta-information representations, and developing other methods for salience manipulation (e.g., image processing algorithms to determine background texture and adjust an entity's texture accordingly.)

3.5 Examples in C2 asset allocation domain

Within the domain of C2 asset allocation, we have developed a set of visualization methods that depict quantities of information and meta-information. One set of requirements identified through the analysis includes the need to understand the role of meta-information affects reasoning about a unit's particular suitability to a task, and how that suitability is affected by environmental factors. One vignette asked the user to consider the suitability of units to the task of photographing river crossings. Information needs included unit flight and ISR capabilities, the threat level of the unknown area, and the current environment. Meta-information needs included the certainty of potential threat location and type, and the certainty and recency of weather reports. Figure 4 displays the visualizations that represent the information and meta-information for this particular situation.

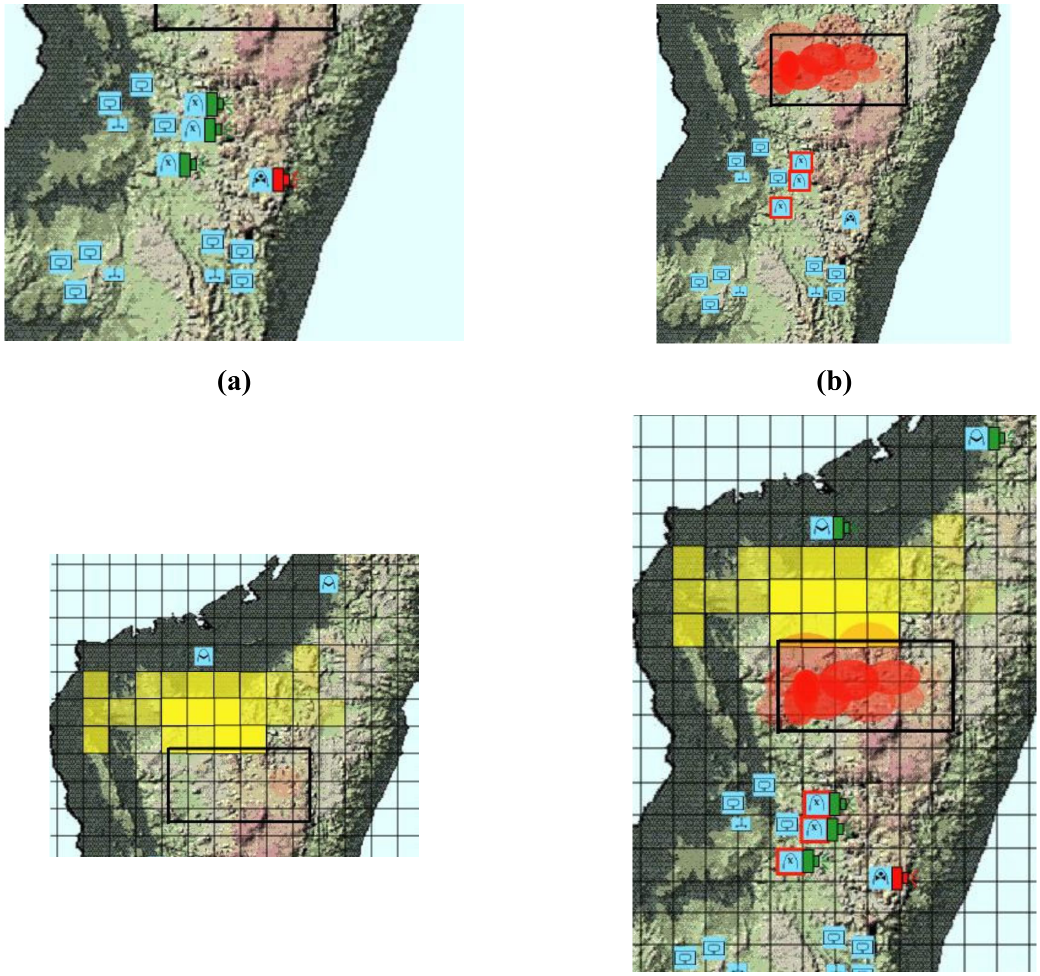

(a) (b)

(c) (d)

Figure 4: Different meta-information representations. (a): Usage of glyphs to indicate unit capabilities. (b): Threat severity representation. (c): Weather report recency representation. (d): Those three methods combined.

In this situation, we experimented with using glyph augmentation to represent unit capabilities (i.e., green or red camera icons indicating the capability or incapability of video capture, and red borders to indicate vulnerable units), and transparency to represent certainty of threat and weather reports.

In another vignette, we consider the problem of reassigning an entity to a task of high priority. Information needs in this situation included current tasking, a unit's total time needed to complete the new task, current payload, threat level in traversed areas, and environmental concerns. Meta-information needs included priority of each current task, factors used to calculate time needed, and each unit's vulnerability to the known threats. Figure 5 displays the visualization developed for this situation. Units that are completely incapable of completing the mission are removed from the display. Transparency is mapped to priority of the current task, with higher priority tasks receiving a higher level of transparency. Icon size is manipulated to indicate the estimated time to task completion. A larger icon indicates that the corresponding unit is more likely to achieve the task in a smaller time frame.

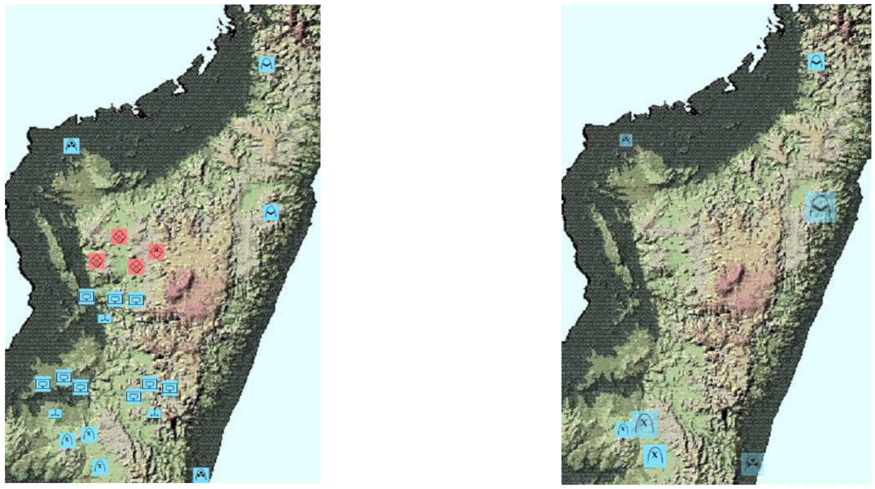

(a) (b)

Figure 5: Meta-information representations for task rescheduling. (a) No meta-information representation; (b) Filtering, transparency, and size used to represent current task priority and estimated time to completion

Our work has produced many visualization techniques that are applicable to many domains. Figure 6 illustrates a series of visualizations across a variety of operating environments and situations.

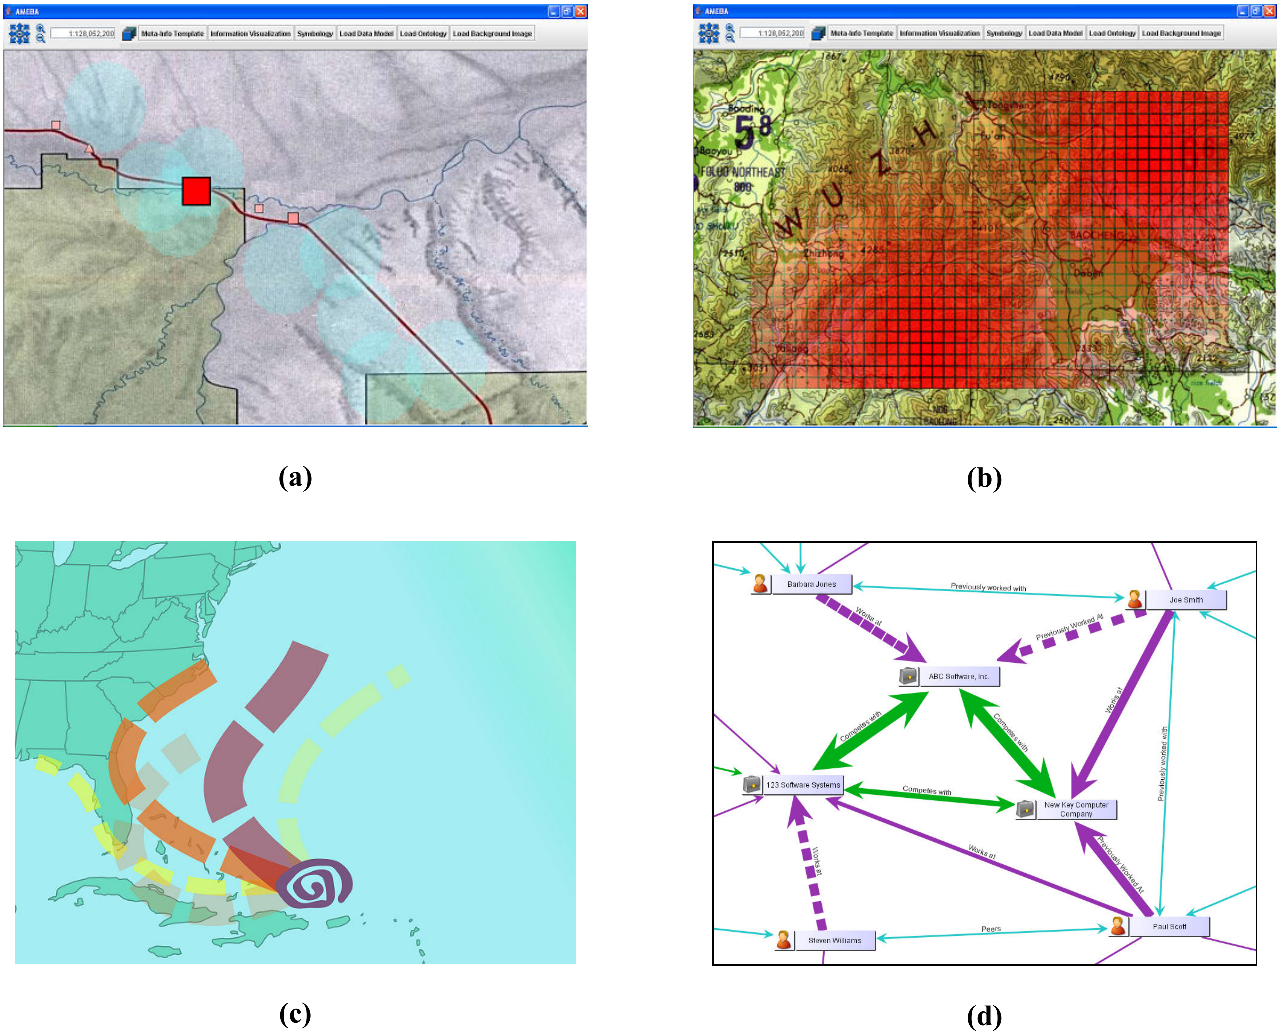

Figure 6: Meta-information representations across several domains. (a): Certainty of unit classification represented through icon size and saturation; (b) Hostile unit range and associated certainty expressed through transparency; (c) Certainty of weather prediction represented through size and hue; (d) Recency and certainty of social network connections represented through dashes and line thickness.

4. Implications and Discussion

The examples presented in this paper highlight the potential impact of meta-information on decision-making and point to implications for future development of C2 support systems. The example methods described here should be understood in the context of prototype system development – the methods are intended to be refined through additional interaction with domain experts and through usability testing. However, these methods do illustrate some of the key features of our approach to providing specific decision-support tools based on an understanding of the role of meta-information in decision-making.

The development of prototype systems to support meta-information has identified a number of issues in the design of displays and interfaces that must be addressed to effectively communicate the needed information and meta-information for a particular task. The addition of meta-information to an already information-rich display clearly may lead to the presentation of too much information in conflicting forms, which could easily occur if meta-information is displayed without thought to how it might optimally compliment the information it is qualifying. The careful holistic design of visualization methods for both the information and meta-information is one approach to prevent overload. It is important to examine the circumstances under which decision-makers will be critically affected by the presence or absence of meta-information. Situations where there are "close calls" or where the risk associated with alternative actions is high are likely to require more careful consideration of meta-information in decision-making to improve the salience of important information while *avoiding* information overload. Another approach is to provide methods for user-selected display of meta-information, although work by Wickens and Yeh (1997) suggests that an effectively designed display of all information is more efficient than a user-controlled display of particular information elements. When designing displays, an effort should be made to determine if there are "natural mappings" (Helander, 1987) between a type of meta-information and a display method that could be exploited to make the presentation of meta-information more intuitive and therefore minimize potential overload problems. Initial results (Bisantz et al., 2006) indicate that for some visual characteristics (e.g., transparency) there are natural orderings for representing meta-information, while for others (e.g., hue) there not. However, more research is needed to determine the impact of information and visual contexts.

In the process of conducting the efforts represented in this paper, we established a number of areas in which future research and development is needed. These include the need to:

- Define clear empirical relationships between human performance and the presence/absence of meta-information
- Further define the specific impacts of meta-information on human perception, reasoning, and action
- Develop principles for visualization and user interface designs that incorporate meta-information
- Understand methods by which humans interact with complex systems through the use of meta-information
- Develop methods for representing complex system internals (and processing their states) to generate needed meta-information

In our ongoing efforts, we hope to begin to address these and other concerns as we develop decision support systems for C2.

5. References

Bisantz, A., Pfautz, J., Stone, R., Roth, E., Thomas-Meyers, G., & Fouse, A. (2006). Assessment of Color Variables for Displaying Meta-Information on Maps. In *Proceedings of Human Factors and Ergonomics Society 50th Annual Meeting*. San Francisco, California.

Durlach, P. (2004). Change Blindness and Its Implications for Complex Monitoring and Control Systems Design and Operator Training. *Human-Computer Interaction*.

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

Endsley, M. R. & Rodgers, M. D. (1994). Situation Awareness Information Requirements for En Route Air Traffic Control. In *Proceedings of Human Factors and Ergonomics Society 38th Annual Meeting*. Santa Monica, CA.

Endsley, M. R. (1993). Situation Awareness and Workload: Flip Sides of the Same Coin. In *Proceedings of Proceedings of the Seventh International Symposium on Avitaion Psychology,* R. S. Jensen & D. Neumeister (Eds.), (pp. 906-911). Columbus, OH: Ohio State University, Department of Aviation.

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

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

Fleetwood, M. D. & Byrne, M. D. (2003). Modeling the Visual Search of Displays: A Revised ACT-R/PM Model of Icon Search Based on Eye Tracking Data. In *Proceedings of Fifth International Conference on Cognitive Modeling,* F. Detje, D. Dorner, & H. Schaub (Eds.), (pp. 87-92). Bamberg, Germany: Universitas-Verlag Bamberg.

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

Guarino, S., Pfautz, J., Cox, Z., & Roth, E. (2006). Modeling Human Reasoning About Meta-Information. In *Proceedings of 4th Bayesian Modeling Applications Workshop at the 22nd Annual Conference on Uncertainty in AI: UAI '06*. Cambridge, Massachusetts.

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

Helander, M. G. (1987). Design of Visual Displays. In G. Salvendy (Ed.), *Handbook of Human Factors*. New York, NY: John Wiley & Sons, Inc.

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

Huang, L. & Pashler, H. (2005). Attention Capacity and Task Difficulty in Visual Search. *Cognition,*(94), B101-B111.

Lefevre, R., Pfautz, J., & Jones, K. (2005). Weather Forecast Uncertainty Management and Display. In *Proceedings of 21st Int'l Conf. on Interactive Information Processing Systems (IIPS) for Meteorology, Oceanography, and Hydrology*. San Diego, CA.

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

Lodha, S. K., Sheehan, R. E., Pang, A. T., & Wittenbrink, C. M. (1996). Visualizing Geometric Uncertainty of Surface Interpolants. In *Proceedings of Graphics Interface 1996,* (pp. 238-245). Canadian Information Processing Society, Toronto ON.

MacEachren, A. M. (1992). Visualizing Uncertain Information. *Cartographic Perspective,* 1310-19.

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.

Norman, D. A. (1988). *The Psychology of Everyday Things*. New York: Basic Books.

Pang, A., Wittenbrink, C., & Lodha, S. (1997a). Approaches to Uncertainty Visualization. *The Visual Computer,* 13(8), 370-390.

Pang, A. T., Wittenbrink, C. M., & Lodha, S. K. (1997b). Approaches to Uncertainty Visualization. *Visual Computing,* 13(8), 370-390.

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

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

Pearl, J. & Russell, S. (2000). *Bayesian Networks.*

Pfautz, J., Bisantz, A., Roth, E., Fouse, A., & Nunes, A. (2006a). Beyond Uncertainty: Examining Meta-Information Visualization Techniques. In *Proceedings of SIGGRAPH*. Boston, MA.

Pfautz, J., Bisantz, A., Roth, E., Fouse, A., & Shuster, K. (2005a). Meta-Information Visualization in Geographic Information Systems. In *Proceedings of Society for Information Display '05*. Boston, MA.

Pfautz, J., Roth, E., Bisantz, A., Fouse, A., Madden, S., & Fichtl, T. (2005b). 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. (2006b). The Role of Meta-Information in C2 Decision-Support Systems. In *Proceedings of Command and Control Research and Technology Symposium*. San Diego, CA.

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. M., & Goodstein, L. P. (1994). *Cognitive Systems Engineering*. New York: Wiley and Sons.

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

Schraagen, J. M. C., Chipman, S. F., & Shalin, V. L. (2000). Introduction to Cognitive Task Analysis. In J. M. C. Schraagen, S. F. Chipman, & V. L. Shalin (Eds.), *Cognitive Task Analysis* (pp. 3-24). Mahwah, NJ: Lawrence Erlbaum Associates.

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

Whitaker, R. & Thomas-Meyers, G. (2006). Knowledge Glyphs: Visualization Theory Development to Support C2 Practice. In *Proceedings of Command and Control Research and Technology Symposium*.

Wickens, C. D. & Yeh, M. (1997). Attentional Filtering and Decluttering Techniques in Battlefield Map Interpretation. In *Proceedings of Advanced Displays and Interactive Displays,* (pp. 2-35). Adelphi, Maryland.

Wolfe, J. M. (1998). Visual Search: A Review. In H. Pashler (Ed.), *Attention*. London, UK: University College London Press.

Woods, D. D., Patterson, E. S., & Roth, E. M. (2002). Can We Ever Escapre From Data Overload? A Cognitive Systems Diagnosis. *Cognition, Technology and Work,* 4(1), 22-36.

Wyszecki, G. & Stiles, W. S. (1982). *Color Science: Concepts and Methods, Quantitative Data and Formulae*. (2nd ed.) New York: John Wiley & Sons.

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