14th ICCRTS
"C2 and Agility"

147: Enabling Robust C2 Systems through Evolvable Human-In-The-Loop Data Fusion

Topic 6: Modeling & Simulation

Samuel Mahoney, Jonathan Pfautz, Ted Fichtl, Sean Guarino, Eric Carlson, Mike Farry
Charles River Analytics, Inc
625 Mount Auburn Street
Cambridge, MA 02138

Emilie Roth
Roth Cognitive Engineering
Brookline, Massachusetts

Gerald Powell
U.S. Army RDECOM CERDEC I2WD
Ft. Monmouth, New Jersey

Enabling Robust C2 Systems through Evolvable Human-In-The-Loop Data Fusion

Abstract:

Data fusion systems are being increasingly used to support military planning, decision making, and command and control functions in general. Typically, these systems are designed around the current capabilities of particular data collectors (e.g., sensors) and available processing algorithms. These algorithms incorporate an "ontology" that reflects the designer's perception of key concepts in the world (e.g., types of threats, classes of vehicles to be tracked) and how these can be parsed by the data fusion systems. As a consequence, these algorithms are limited in their ability to adapt to the dynamic changes that inevitably arise in the operational environment (e.g., new sensors, weapons, and enemy tactics). This frailty is representative of a more generic problem with current approaches to system design that result in rigid systems that are unable to evolve to keep pace with changing operational conditions. In this paper, we present the results of an analysis, design, and development effort intended to move towards robust C2 through evolvable human-in-the-loop data fusion systems. We discuss an evolvable semantic interface we have designed that enables the creation of new concepts within the fusion system, and provide an overview of the prototype evolvable data fusion system architecture we are developing.

(199 Words)

1 INTRODUCTION

Data fusion systems are being increasingly used to support military planning, decision making, and command and control functions in general. These systems combine, correlate, and aggregate heterogeneous and distributed sources of information with the goal of providing needed information (Waltz & Llinas, 1990) – for example, combining the inputs from multiple ground-based radar systems to track an adversary convoy.

Data fusion systems typically consist of data structures that are commonly termed "ontologies" (e.g., this precipitation sensor reports one of one of the following: rain, sleet, hail, snow) and inference mechanisms (e.g., a rule states "if sensor X reports **`wet`** and sensor Y reports **`10°C`** then report **`raining`**"). These ontologies and other representational formalisms are typically specified as part of the initial design of the data fusion system. They reflect capabilities of particular data collectors (e.g., sensors) and processing algorithms (e.g., correlating target locations over time) at the time the system was initially developed, as well as the designer's perception of the key features of the world (e.g., types of threats, classes of vehicles to be tracked) and how these can be parsed by the data fusion systems.

A consequence of this design approach is that it results in data fusion systems that structure information according to the system's ability to perceive the world, rather than according to a human's need to understand and act upon the world. Further, these data fusion systems are limited in their ability to adapt to dynamic changes that inevitably arise in the operational environment. Inevitably, there are changes in both "own" and “adversary” assets (e.g., sensors, weapons, equipment) and operations that cannot be accommodated by rigidly designed data fusion systems.

This frailty is representative of a more generic problem with current approaches to system design that result in rigid systems that are unable to adapt to rapidly changing operational environments (Pew & Mavor, 2007; Roth et al., 2006). In this paper, we present the ongoing results of an analysis, design, and development effort intended to move away from traditional data fusion systems towards evolvable human-in-the-loop data fusion systems.

Our research and development is being performed as part of an Army program examining data fusion methods that support rapid knowledge building and editing, enabling data fusions systems that will be knowledge-intensive and will respond to a changing battlefield environment. For example, fusion systems must cope with new threat doctrine, varying Tactics, Techniques, and Procedures (TTPs), and equipment or weapon changes. A key goal of the program is to develop practical, operational tools and systems.

In support of this program, we have adopted a cognitive systems engineering approach targeting the design of evolvable support (Roth et al., 2006) for data fusion systems. In this paper, we present some relevant background material used to motivate our work, discuss the methods used to perform an analysis in support of an evolvable system design, and provide an overview of the prototype evolvable data fusion system architecture we are developing.

2 BACKGROUND

2.1 Human-In-The-Loop Data Fusion

Fusion systems generally provide C2 with valuable information aggregated from multiple information sources through computation reasoning methods. More specifically, fusion algorithms are categorized in a stratification depending on their purpose (Steinberg, Bowman, & White, 1998; U.S.Department of Defense, 1991). For example, lower level fusion algorithms reconcile sensor reports to specific entities and aggregate these reports into entity tracks. In higher level fusion algorithms, relationships and impacts are identified and predicted based on tracks established and maintained at the lower levels. Popular algorithms used in various level of reasoning include Kalman filtering (Kohler, 1997; Kalman, 1960), joint probabilistic modeling (Ahmeda et al., 1997), and particle filtering (Das et al., 2005; Koichiro, Kawanaka, & Okatani, 2004; Gordon, Simon, & Kirubarajan, 2002), all of which employ various stochastic and probabilistic hypothesis generation and selection methods to reason about data ingested from available sensors and systems.

There is clear recognition in the data fusion community that the role of the human in the data fusion process could be exploited to a far greater degree than in current systems, and that there are performance improvements to be gained from including human knowledge in the process. For example, Blasch and Plano (2003, 2002) have argued for a need to adopt a human-in-the-loop approach to data fusion (i.e., the need for a “Level 5” in the JDL Data Fusion Model). Similarly, others have recommended improvements to data fusion to incorporate better mechanisms for supporting analyst understanding of the process and addressing issues related to the qualifiers of information, or *meta-*

information, i.e., quality control, pedigree, reliability, and consistency (Pfautz et al., 2007; Llinas et al., 2004).

Too often, however, system designers and developers solely focus on internal concepts that support primary interoperability among systems and computational function (e.g., fusion algorithms) and do not focus on representational formalisms that support human reasoning. The ability to exploit human capabilities as part of military fusion systems requires an approach that can translate significant algorithmic complexity into a human-accessible concept representation, and one that can translate human input in that representation back into specific impacts on the underlying algorithms. Ontologies can provide the foundation of such an approach.

2.2 Defining Fusion System Concepts

Ontologies explicitly define 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 a research area with wide applications in knowledge-based intelligent systems. Ontologies are critical in formalizing statements in an application domain and in operating with the associated semantics of the concepts (i.e., to provide domain-relevant structures upon which computational methods can act). For this very reason, ontologies can offer strong support not only for building knowledge bases for computational data fusion systems, but also for describing the contexts in which the knowledge is needed by the system.

However, there is an inherent limitation in ontologies defined at the time of data fusion system design in that they can only be used to comprehend terms and concepts that have been pre-identified. This places severe restrictions on their application within dynamic, emerging environments found in military operations, where the core set of concepts and lexicon is constantly changing (e.g., constantly changing information sources, technology, tactics, and organizational structures of both own and adversary forces).

While there have been developments to ease the authoring of ontologies (Tablan et al., 2006; Farquhar, Fikes, & Rice, 1997), we identified a need to design software structures and user interfaces that would allow for the creation and adaptation of ontologies as part of the operational system. This capability would allow the system to evolve to accommodate changes in the operational environment. The nature and form of the evolution should necessarily be driven by an analysis of the work domain that identifies potential weak points in a representation (e.g., we know that weather data can be reasonably well-defined in a static way, but that adversary use of communications technologies varies dramatically from month to month).

2.3 A Need for Evolvable Systems

There is growing recognition that the activities that people engage in and the physical, social, and organizational environment in which these activities take place are constantly evolving (Roth et al., 2006; Woods & Dekker, 2000). Operational military intelligence personnel face consistent revisions to the goals, scale, scope, structure, and information sources entailed by their job function (Roth et al., 2006). However, the technological systems that they interact with are built with predefined static structures that cannot be easily modified to keep pace with the changing conditions (Truex, Baskerville, & Klein, 1999). As the work domain inevitably evolves, users are often forced to devise workaround solutions which combine internal system elements and external technology.

There is a growing call to develop efficient techniques that can dynamically capture changes in both work context and requirements and to also create "evolvable" systems that can be readily adapted to meet changing conditions of work (e.g., (Pew et al., 2007; Roth et al., 2006; Hoffman & Elm, 2006)). The goal is to ease additional workload and collaborative discontinuity that workarounds may cause by anticipating particularly vulnerable aspects of a system to operational changes (Roth et al., 2006). We adopted this approach towards development of an evolvable human-in-the-loop data fusion system to address concerns with overly static (and hence "brittle") data fusion systems, but also to understand the practical system engineering challenges inherent in such an approach.

3 ANALYSIS AND DESIGN OF AN EVOLVABLE DATA FUSION SYSTEM

We are developing a prototype human-in-the-loop data fusion system as part of an Army research program. Our focus is targeting a data fusion system that supports the assessment of different operational strategies or Courses of Action (COAs). The user of the system can enter one or more alternative COAs and have the system provide an assessment of the likelihood of success of that COA. Our prototype system guides data fusion processes by allowing the user to describe both their current operational goals and the background operational environment or situation. This description then informs underlying data fusion processes to guide the collection, correlation, and aggregation of information that is better tailored to the implicit and explicit needs of the user.

In the sections below, we describe the cognitive engineering processes we employed in designing and developing the system. After an initial domain analysis, we began by developing an initial prototype that relied on a predefined ontology. We quickly came to realize that we needed to include mechanisms to enable the users of the system to extend the ontology so as to be able to cope with an ever changing operational environment; this led to a second cycle of design that focused on incorporating evolvable features in the prototype. We describe the core data fusion system we developed and the features we incorporated into the system to allow the ontology to be extendable by the user community.

3.1 Cognitive Analysis

We performed an initial cognitive task analysis intended to provide a broad characterization of the Military Intelligence (MI) domain and the sources of cognitive demands and performance challenges. Analysis was based on a series of knowledge elicitation sessions conducted with three Subject Matter Experts (SMEs), all of whom are former Army Military Intelligence officers with extensive intelligence analysis experience. This focused analysis was then validated and extended based on subsequent interviews with current military intelligence personnel and prototype feedback evaluation sessions. These efforts represent over 600 hours of interviews and evaluations with over 45 different subject matter experts. Additionally, these efforts included discussion of historic, current, and planned data fusion systems used in Army intelligence analysis (e.g., ASAS, ASAS-Lite, DCGS-A).

A more focused analysis was then conducted to define the ontologies and associated representational formalisms that would support human-in-the-loop data fusion. Our prototype system guides data fusion processes by allowing users to describe both their current operational goals and the operational environment. This description then informs underlying data fusion processes to guide the collection, correlation, and aggregation of information that is better tailored to the implicit and explicit needs of the user. For example, we identified that a key, often un-articulated aspect of an operational situation is the degree to which air-based assets will be used. This aspect, often assumed by the user to be readily apparent, is not typically used in guiding data fusion, but has a clear impact on the degree of processing required (i.e., it modifies the number of potential threats identified and tracked to include any ground-to-air or air-to-air threats). As a consequence, we determined that it was important for our data fusion system to capture not only the user's operational goals but also the general operational situation that constitutes the background context.

The initial challenge was to develop a process that allows the system users to express their understanding of operational goals and operational situations. An initial set of operational goals and operational situation descriptions was collected via structured interviews with a core set of two SMEs. The interviews focused on description of actual past cases as well as analysis of responses to simulated cases exercises. From these interviews, we derived an initial set of questions that could be used by the data fusion system to characterize operational goals and situation descriptions.

The next step in our analysis was to work from these potential questions to the space of possible responses. While some cases were simple (e.g., in the above example, a "yes/no" response was expected), others were more complicated, and a set of branching query statements were identified (e.g., "if you know it is a wheeled vehicle, then you need to ask if its speed and heading are indicative of this type of threat"). The set of possible responses captured in an ontology needed to be developed at a level of abstraction that supported relatively rapid response to the queries, while also containing enough detail to provide a significant impact on the data fusion process. This presented a particular challenge in our analysis, and we found that multiple iterations of interviews and prototype evaluations where needed to identify an effective level of abstraction of key operational concepts and operational goals.

The initial conceptual framework was exercised and refined using a corpus of representative operational goals and operational situation descriptions collected from a broader range of operational users. In total, we collected over 100 representative operational goal descriptions via interviews and exercises using an initial system prototype. This corpus of representative goals and situation descriptions was used to iteratively develop and test the set of questions to be used by the system to elicit and characterize the user's operational goal and background operational situation.

3.2 System Overview

Figure 1 depicts the functional flow of our prototype. To perform a given *COA impact assessment*, the system elicits current *operational goals* and *operational situations* from the user through an iterative *questioning process*. The prototype then calculates the potential performance of a user-selected COA against *dimensions of performance* derived from the operational goals and operational situation as defined by the user. To facilitate this calculation, we developed a SME populated database of course of action elements and their characterization across the "universe" of defined dimensions of performance.

For example, our prototype will help assess how effective one can expect a selected COA to proceed when the enemy employs asymmetric tactics (e.g., covert munitions, non-military communications) vs. symmetric tactics (e.g., overt munitions, military communications). Our prototype can then contrast expected impact if the enemy chooses to operate in an urban environment rather than a rural one.

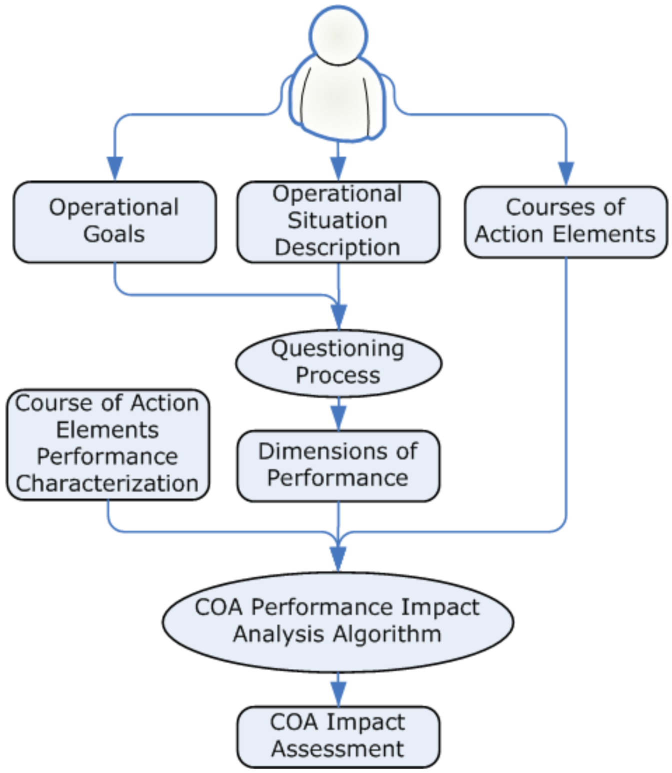

Figure 1: Overview of prototype impact assessment human-in-the-loop fusion process

3.3 Preliminary Approaches

Our prototype has undergone a series of design and test cycles that have involved user evaluations of functioning software prototypes. These evaluations used military personnel with current operational experience as test participants. They exercised the prototype using operational scenarios (both ones we developed and ones they provided). These evaluations provided an opportunity to obtain user feedback as well as to expand our corpus of cases to use in system development and test.

Our preliminary attempts to define the user-facing system formalisms resulted in a fixed ontology supporting user interactions with the system. This fixed ontology represented answers to the questioning process shown in Figure 1. For example, to describe enemy assets, we provided the user a list of weapon categories similar to that of the Military Scenario Definition Language (MSDL; http://www.sisostds.org), such as `large crew-served weapon` or `covert hand-held weapon.` Figure 2 shows a simplified representation our initial approach.

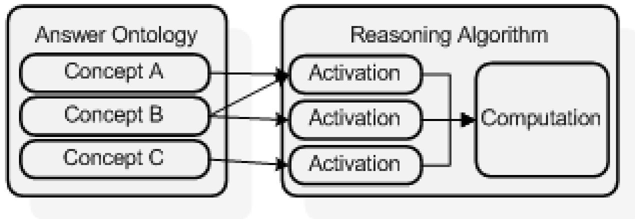

Figure 2: Fixed Ontology approach to HITL Fusion

In this prototype iteration, the users selected an answer from the ontology which evokes a numerical activation within the reasoning algorithm. These activations are stored within the ontology as attributes of the concepts. We evaluated the prototype with Army SMEs who were able to quickly develop operational contexts that they had experienced but were not expressible with our fixed ontology. We subsequently made addendums to the ontology accounting for a move to Stability and Support Operations (SASO) from High Intensity Conflict (HIC). Again, the operational environment had shifted predominantly to a Counter Insurgency (COIN) Campaign, and our ontology failed to capture a majority of the operational situations invented by the SMEs.

These ontological failures were not just the result of missing concepts. In some cases, there are mismatches in the terminology that represent the concepts. In other cases, the semantics of the operational concepts drift as an operation unfolds, which would result in updated activations within the fusion reasoning algorithm (e.g., “What did an IED entail in 2004 versus today?”). After a few iterations of design, implementation, and testing by Army SMEs that recently returned from OIF, it became clear that the evolving operational contexts could never be confined to a fixed ontology.

3.4 Incorporating Evolvable Design Features

Feedback from the preliminary evaluations made it clear that capabilities are needed to allow users to expand the pre-defined corpus of terminology by which a user communicates the operational goals and situation. We found great variability in the terms used by participants from different operational environments. We also found that the terms and concepts used to define an operational goal and operational situation could change over time. In particular, new doctrine for targeting emerging adversary tactics was released at multiple points through our study. In addition, our analysis and evaluations revealed a consistent desire on the part of our evaluation participants to use either the most comfortable or most up-to-date terminology to describe operational goals and operational situations. We concluded that we needed to develop evolvable components to our system to reflect these ever-changing system requirements.

To accomplish this goal, we needed to define not only the goal-based and situational impacts on our data fusion process, but also the dimensions of these impacts that would allow a user to extend or refine the underlying model while still maintaining system function. This user-centered meta-structure was developed through the same iterative interview and evaluation process. Figure 3 shows an updated abstraction of our current system.

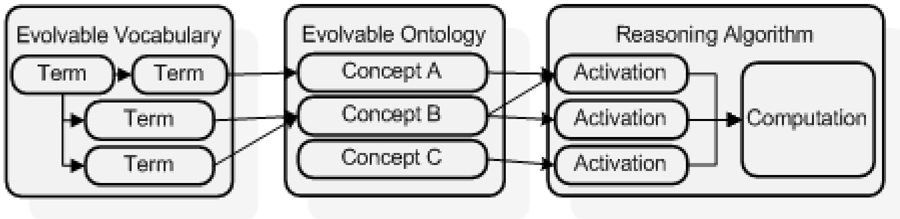

Figure 3: Evolvable HITL Fusion Prototype

When users encounter a situation where they are unable to express operational concepts with existing vocabulary, they can define a new term (e.g., "the adversary is using a new type of vehicle I'm going to call an **`armored scout vehicle`** that has the following properties…"). This requires the user to either map the new term to concepts in the existing vocabulary (e.g., "an **`armored scout vehicle`** is going to be the same as a **`pickup truck`** in terms of a COA"), or add additional concepts and characterize their impact across the various dimensions of performance (e.g., "an **`armored scout vehicle`** is going to have a particular effect that is unique and will alter the predicted performance of the COA").

In this way, our underlying system ontology can grow to accommodate dynamic operational changes. For example, if the system initially recognized the concept **`armored tracked vehicles`**, and then adversary tactics changed to using faster, lightly armored vehicles, the user should be able to create a new class of operational targets, specifying that one (among many) differentiator in identifying such targets would be the difference in the speed of the target object. This flexibility allows the user the freedom to define terminology that may be local to the unit, and define in terms that are semantically valid to the fusion algorithm. Further, the process of mapping each term to concepts captured within the ontology can solidify the user's understanding of which aspects of a given entity can influence the fusion algorithm. A screen shot of our *Term Editor* interface is shown in Figure 4.

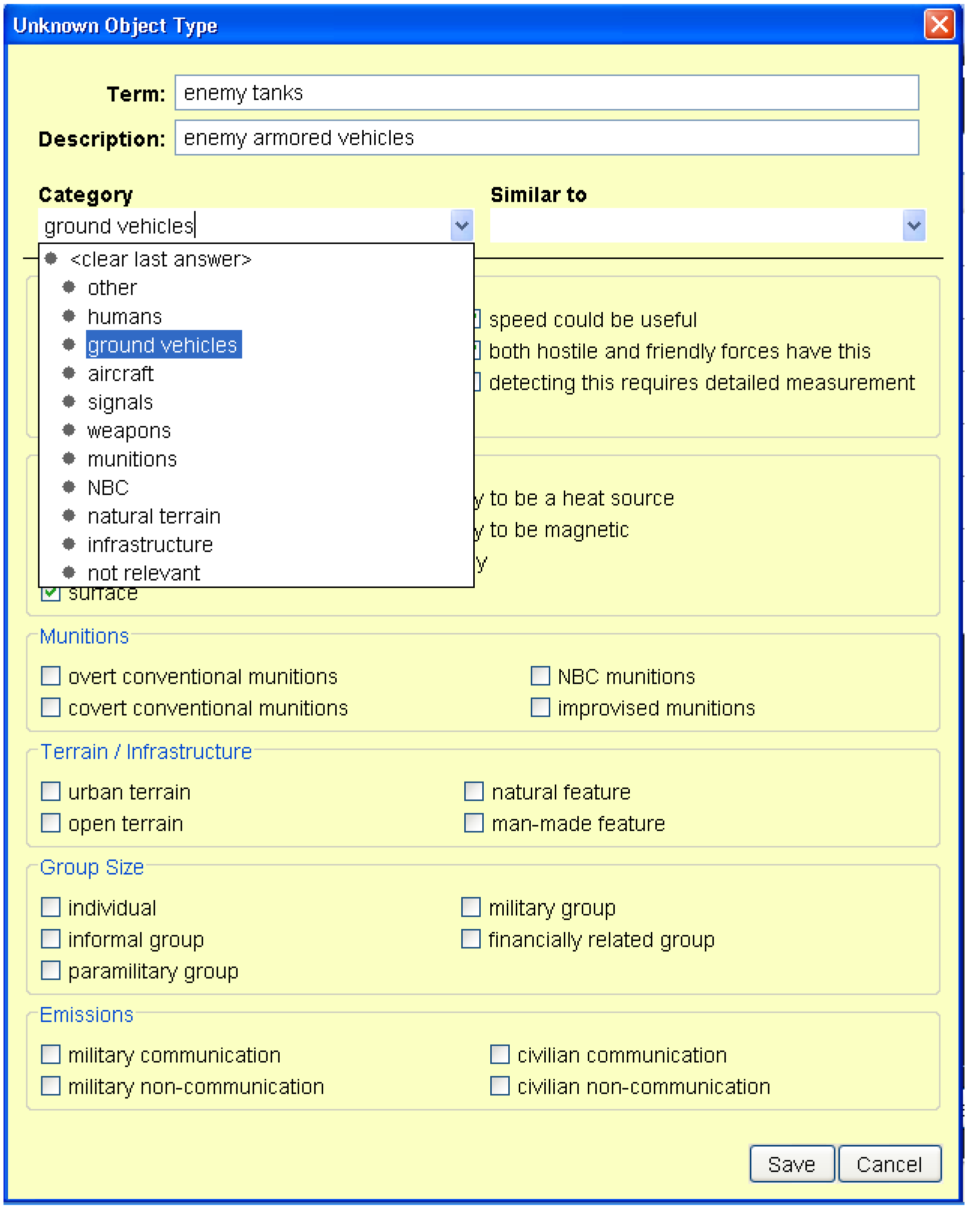

Figure 4: Term Editor Interface

In this screen shot, the user has entered a term, `enemy tanks`, which the prototype does not currently recognize. The user can then categorize the term within the existing vocabulary hierarchy, or declare it synonymous with an existing term. In either case, the system populates the properties of the associated term that the user can then alter. These properties link the term with the ontology that influences the fusion algorithm.

If the concepts defined within the system ontology do not accurately capture the essence of an undefined term, a separate process is available for adding new concepts and defining influences within the COA algorithm. This process is lengthy in comparison to the addition to new terminology, and is expected to be necessary less frequently. Newly added concepts are then available for mapping to existing and newly defined terms. By developing these interfaces that enable the user to expand or alter the concepts within the ontology, we ensure that the system can continue to function without engineering intervention.

Finally, the underlying fusion algorithm must incorporate several properties to enable these evolvable capabilities. First, the algorithm must accept new elements feeding the computation. Second, there must be a means for enumerating the influences of these sources. This can be as simple as accepting new states and probabilities, or as complex as defining entire causal influence models (Pfautz et al., 2009; Cox & Pfautz, 2007). Third, the algorithm must accept and validly handle a means for defining uncertainty. In other words, the user needs the ability to say "I don't know" in the face of an unanticipated case not well supported by abstraction.

4 IMPLICATIONS AND FUTURE WORK

Our work is intended to be a practical application of the principles of evolvable work centered design. We identified the need for users to expand the set of terms the data fusion prototype is able to understand and reason about and developed facilities to enable this need to be fulfilled. In this way, the prototype enables the underlying ontology to grow and evolve to keep pace with dynamic changes in the operational environment. This approach will be valuable as standard concept models, such as the Joint Command, Control and Consultation Information Exchange Data Model (JC3IEDM) further facilitate intra- and inter-system reasoning. We are currently undertaking additional efforts to assess the interoperability impacts of evolving user-facing ontologies. We plan to conduct further evaluations over the next year to assess whether the software truly provides the flexibility that is intended. Clearly, the expectation is that the system vocabulary will be more volatile than evolving ontology. As part of our evaluations, we will capture statistics regarding user to user differences in vocabulary and ontology updates.

In our current implementation, each user's model remains local, as it is a personalized representation of the work domain. We plan to develop methods for aggregating and analyzing individual ontologies (manually, semi-automatically, or automatically) to establish updated universal baselines which can be used in multiple work domains. A resulting benefit of this approach is that it may lead to the identification of inconsistencies in results between users for the same work domain and therefore aid in empirically locating cognitive inconsistencies between team members and identifying misconceptions about the work domain. That is, an evolvable system has the potential to provide feedback not only to improve itself, but to aid the designers of the system in providing revisions.

5 ACKNOWLEGEMENTS

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 remain anonymous, we would like to particularly thank MAJ David Lorenz (ret.) and MAJ Chet Brown (ret.) for their contributions.

6 REFERENCES

Ahmeda, S. S., Keche, M., Harrison, I., & Woolfson, M. S. (1997). Adaptive Joint Probabilistic Data Association Algorithm for Tracking Multiple Targets in Cluttered Environment. Radar, Sonar and Navigation, IEEE Proceedings 144[6].
Ref Type: Generic

Cox, Z. & Pfautz, J. (2007). *Causal Influence Models: A Method for Simplifying Construction of Bayesian Networks.* (Rep. No. R-BN07-01). Cambridge, MA: Charles River Analytics Inc.

Das, S., Lawless, D., Ng, B., & Pfeffer, A. (2005). Factored Particle Filtering for Situation Assessment in Urban Environments. In *Proceedings of 8th International Conference on Information Fusion*. Philadelphia, PA.

Dujmovic, J. J. & Larsen, H. L. (2007). Generalized Conjunction/Disjunction. *International Journal of Approximate Reasoning,* 46(3), 423-446.

Farquhar, A., Fikes, R., & Rice, J. (1997). Tools for Assembling Modular Ontologies in Ontolingua. In *Proceedings of AAAI*.

Gordon, N., Simon, M., & Kirubarajan, T. (2002). Efficient Particle Filters for Joint Tracking and Classification. In *Proceedings of Signal and Data Processing of Small Targets*. Orlando, FL: SPIE.

Hoffman, R. & Elm, W. C. (2006). HCC Implications for the Procurement Process. *IEEE Intelligent Systems,* 21(1), 74-81.

Kalman, R. E. (1960). A New Approach to Linear Filtering and Prediction Problems. *Transactions of the ASME,*35-45.

Kohler, M. (1997). *Using the Kalman Filter to Track Human Interactive Motion -- Modelling and Initialization of the Kalman Filter for Translational Motion.* (Rep. No. 629). University of Dortmund.

Koichiro, D., Kawanaka, O., & Okatani, T. (2004). Object Tracking by the Mean-Shift of Regional Color Distribution Combined With the Particle-Filter Algorithm.

Llinas, J., Bowman, C., Rogova, G., Steinberg, A., Waltz, E., & White, F. (2004). Revisiting the JDL Data Fusion Model II. In *Proceedings of 7th Int. Conf. on Information Fusion*.

Pew, R. W. & Mavor, A. S. (2007). *Human-System Integration in the System Development Process: A New Look*. Washington, DC: Committee on Human Factors, Division of Behavioral and Social Sciences and Education, National Research Council, National Academies Press.

Pfautz, J., Carlson, E., Koelle, D., Roth, E., & Cox, Z. (2009). Rapid Creation and Validation of Bayesian Belief Networks With Causal Influence Models. *International Journal of Approximate Reasoning,* In print.

Pfautz, J., Fouse, A., Farry, M., Bisantz, A., & Roth, E. (2007). Representing Meta-Information to Support C2 Decision Making. In *Proceedings of International Command and Control Research and Technology Symposium (ICCRTS).*

Quine, W. V. (1969). *Ontological Relativity and Other Essays*. New York: Columbia University Press.

Roth, E., Scott, R., Deutsch, S., Kuper, S., Schmidt, V., Stilson, M. et al. (2006). Evolvable Work-Centered Support Systems for Command and Control: Creating Systems Users Can Adapt to Meet Changing Demands. *Ergonomics,* 49(7), 688-705.

Steinberg, A., Bowman, C., & White, F. (1998). Revisions to the JDL Data Fusion Model. In *Proceedings of Joint NATO/IRIS Conference*. Quebec, Canada.

Tablan, V., Polajnar, T., Cunningham, H., & Bontcheva, K. (2006). User-Friendly Ontology Authoring Using a Controlled Language. In *Proceedings of 5th International Conference on Language*.

Truex, D., Baskerville, R., & Klein, H. (1999). Growing Systems in an Emergent Organization. *Communications ACM*, 42117-123.

U.S.Department of Defense, D. F. S. o. t. J. D. o. L. T. P. f. C. (1991). *Data Fusion Lexicon.*

Waltz, E. & Llinas, J. (1990). *Multisensor Data Fusion*. Norwood, MA: Artech House, Inc.

Woods, D. & Dekker, S. (2000). Anticipating the Effects of Technological Change: A New Era of Dynamics for Human Factors. *Theoretical Issues in Ergonomic Science,* 1(3), 272-282.