
Volume 10, issue 1, January 2009
ISSN 1566-2535
ELSEVIER
An International Journal on
INFORMATION FUSION
Special Issue on
High-level Information Fusion
Multi-Sensor, Multi-Source
Information Fusion
Editor-in-Chief
Belur V. Dasarathy
Available online at www.sciencedirect.com
ScienceDirect
Now indexed by Thomson Scientific

Information Fusion 10 (2009) 2–5

Contents lists available at ScienceDirect

Information Fusion

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

Guest Editorial

High-level information fusion and situation awareness

High-level information fusion includes situation assessment, impact assessment and process refinement. Within the JDL Data Fusion Model [1], these aspects of fusion fall in levels 2, 3 and 4. Under situation assessment, the goals include deriving higher-order relations and identifying meaningful events and activities. Impact assessment includes estimating the level of threat or danger, predicting possible outcomes of particular decisions, determining the vulnerabilities of ones own assets and determining possible courses of actions. Refinement of the fusion process to improve the fusion from level 0 through level 3 represents the next critical stage. Fusion at this level also includes dynamic sensor management.

In view of the growing importance of these fusion objectives, the journal Information Fusion had solicited the publication of a special issue devoted specifically to high-level information fusion and situation awareness. The aim of this special issue was to provide a focal point for recent advances in the area of high-level information fusion across different paradigms and disciplines. Submitters were asked to report new contributions underpinning fusion of information in this domain. In particular, the following list of topics covers some of the issues that were sought for this special issue:

- Situation and threat assessment
- Situation development
- High-level Information fusion test beds
- Multi-sensor, multi-source fusion system architectures for situation awareness
- Battlespace awareness
- Computational intelligence techniques for high-level information fusion
- Common operational picture and intelligence fusion
- Information fusion design and methodology
- Decision fusion
- Context-based fusion and aggregation processes
- Conflict management in high-level fusion
- Knowledge based and probabilistic reasoning in high-level fusion
- Information pedigrees
- Resource management algorithms
- Ontology-based approaches for high-level information fusion
- Statistical and probabilistic reasoning for high-level fusion
- Advanced analysis and intent inference
- Belief analysis
- Intelligent software agents for high-level information fusion
- Real-world applications of high-level information fusion

As the extensive list of topics listed above indicates, high-level information fusion is a very complex domain. For roughly 10 years now the research community has been recognizing the need for significant progress in this domain, calling for new approaches to the problems posed by the real world applications. Sources of solutions have been sought primarily in two domains:

- Levels 0 and 1 data fusion,
- Artificial Intelligence, Computer Science.

The traditional level 0 and 1 data fusion paradigm brings the methods for combining multiple (mainly homogenous) data sources into one fused outcome. The methods of level 0 deal with signals, while level 1 deals with objects. Both levels rely on the knowledge of the underlying model of the object (e.g., the motion model). The processing thus is based on a prescribed (fixed) sequence of mathematical operations for estimation of the features associated with the objects. The applicability of these methods to level 2 and higher is limited due to various reasons:

- The need to deal with relations, rather than with properties of objects.
- The lack of the underlying quantitative model for the entities that are of interest to level 2 (i.e., relations). For instance, the fact that the relation `calls-for-help(driver, service-station)` holds at t, does not imply that the probability of this relation to hold at $t+1$ is high. In some cases, it can be true, but in some others not true.
- The high number of relations that might need to be considered. While in level 1 relations are limited to only spatio-temporal relations, in higher levels all kind of relations may be relevant. It is practically impossible to provide a list that is any close to be exhaustive.
- The lack of a "relevance theory" for relations in higher levels. If there are so many relations to consider, how can we decide which ones to focus on?
- The inability to predict all possible control structures in processing of incoming information. Since it is not known at design time which of the relations may become relevant and in what order, it is not possible to prescribe a sequence in which the various algorithms would be invoked; the invocation sequence, instead, needs to be decided at run time.

The search for solutions to these problems lead many researchers to consider various methods used in Computer Science and Artificial Intelligence. The first problem – dealing with relations –

1566-2535/$ - see front matter
doi:10.1016/j.inffus.2008.05.003

leads to methods that would allow for inferring whether a relation holds or not. This is because relations don't manifest themselves by providing a measurable value (''holds'' or ''does not hold''). The lack of an underlying quantitative model implies that either the search for relations needs to explore a wider space or some other types of models need to be considered (e.g., Bayesian network models that capture dependencies among various objects and object properties). The large number of possible relations implies the need for some knowledge that would limit the number of possibilities, e.g., by explicitly specifying which types of relations are of interest or by specifying an ontology. And finally, the inability to determine the control structure of a level 2 processing system leads to the use of declarative approaches in which ''chunks'' of knowledge are encoded as rules and then an application dynamically selects the rules whenever the preconditions on their applicability are satisfied.

This special issue presents seven papers on issues of high-level information fusion in general and on situation awareness in particular. The topics addressed by these papers range from proposals of universal frameworks that offer to solve all the issues in high-level information fusion to algorithms that solve smaller bounded high-level information fusion problems. Although the classification into these two groups is not crisp, we order the papers in this issue essentially by their affinity to these two types. Consequently, the first three papers *Lambert, Sycara et al., McMichael et al.* are classified in the first group and the papers *Pfeffer et al.* and *Yang et al.* are classified in the second group. The papers *Little and Rogova* and *Kokar et al.* offer relatively general ontologies, yet they do not propose a universal framework. Consequently, these two papers are positioned between the first and the second group.

The first paper (*A Blueprint for Higher-Level Fusion Systems*, by Lambert) addresses a number of issues in information fusion in general. First of all it provides an explanation of why we should start dealing with *relations* rather than just with objects as the ''aspects of interest in the environment''. Referring to his earlier papers, the author discusses why state vectors (commonly used in level 1 fusion) are less expressive than relations (e.g., one cannot say 'something is targeting something else', since there is no object to which such a state vector could be attached). Consequently, the author argues that higher-level fusion requires methods for dealing with symbols in a much deeper sense than level 1. Moreover, structure of concepts is needed in order to ''register'' facts. This is to replace the physical coordinate registration used in level 1.

The second major focus of the paper is an explicit inclusion of the human in the distributed information fusion process. Towards this aim, the author ''deconstructs'' the JDL model so that human, machine and integration *interpretations* can be explicitly considered. This allows the author to discuss various relations between the levels of the model and the interpretations.

The third aspect of this paper is the author's work on the realization of the underlying philosophical ideas in a ''semantic machine'' - ATTITUDE. ATTITUDE is a cognitive machine, meaning that its components include a short time memory and a long time memory with some attitudes, like 'believes' or 'desires' and some routine cognitive behaviors. The cognitive routines are represented as state transition machines. Beliefs are represented as Horn clauses, where *believes* appears as an attitude in some of the clauses. State transition diagrams for Object, Situation and Impact assessment are presented. The estimation process captured in these machines is similar to object detection and tracking in level 1, following the Prediction–Observation–Explanation pattern. The author has experimented with expressing knowledge both in terms of an ontology (in description logics) and in a first-order logic.

The second paper (*An Integrated Approach to High-Level Information Fusion*, by Sycara et al.) presents another attempt to provide a computational architecture to cover all levels of information fusion as classified by the JDL model. The basic tasks described in this paper include force structure recognition, terrain analysis, intent inference and sensor management.

Force structure recognition deals with the aggregation problem, e.g., assigning tanks to platoons. It first performs object recognition, which is followed by clustering and then recognition of the force by comparing the clusters to doctrinal templates.

Terrain analysis identifies terrain features like obstacles, mobility corridors, engagement areas, avenues of approach. The first step of this process, trafficability analysis, identifies areas of interest and classifies them as GO, SLOW-GO or NO-GO for a given type of vehicle. Then the areas are assessed with respect to the trafficability for a given type of unit (canalizing) and the topology of mobility corridors is calculated. The results of this process serve as input to the next processing stage – computation of engagement areas, avenues of approach, observation points and sensor management. An interesting analogy to circuit theory is used in the above analysis.

One of the issues addressed in this paper is the incorporation of context into the fusion process. The work has been influenced by the U.S. Army's Intelligence Preparation of the Battlefield (IPB) doctrine and thus it relies on the notion of context as defined in the Army field manuals where context is classified into six categories called METT-TC (Mission, Enemy, Terrain, Troops available, Time, and Civilian considerations). The incorporation of context is achieved by providing computational processes to compute various METT-TC related features and then using them in solving the fusion tasks. The treatment of uncertainty uses the Dempster-Shafer approach.

Intent Inference computes specific intents, e.g., 'attack', 'occupy'. The analogy of potential field, also used in robotics to link low level controls with high level goals, is utilized here. Intents are represented by units' goals (geo locations) that are associated with realizations of particular intents. Then the paths to these goal locations are analyzed with respect to the resistance along each path. This is then used to compare hypotheses about enemy intents.

Finally, an asset (sensor) management algorithm is presented. This part is classified in the paper as JDL's process improvement. Coordination among sensors is related to which information should be passed to a neighbor. It is based on reinforcements. All types of information (tokens) need to be known and a ''relevance'' matrix needs to be defined upfront.

The system was tested using a military simulation testbed (OTBSAF). Moreover, the system was validated by Subject Matter Experts, who were satisfied with the results.

In the third paper (*Force Deployment Analysis with Generalized Grammar*, by McMichael et al.), the notion of situation assessment is extended to include not only the assessment of the current relations among level 1 situation objects and projection into the future, but also the history of such relations. The paper focuses on finding interpretations of histories represented as parse trees.

While the intent is to provide a generic approach to this kind of situation assessment, the focus of this paper is on force deployment assessment, considered here as a core area of situation assessment. In other words, situation assessment is the generation and assessment of hypotheses about history - activities, groupings and interactions of the force components. The relations in this case are spatio-temporal relations that are expected to be relevant ahead of the processing.

The paper proposes a ''grammatical approach'' that combines ''syntactic inference with semantic representation''. Syntactic inference captures the manipulation of data, while the semantic processing captures the association step, e.g., the association of contacts with tracks in level 1 processing. The major issue is the syntactic processing, since the semantic processing in this case is accomplished by a set of rules that are simpler than the grammar rules. The semantics is expressed in terms of predicate logic expressions. In logic, this kind of semantics would be termed "axiomatic semantics".

The authors come to the conclusion that for level 2 processing a context free grammar is needed. In this direction, they provide a grammar – Generalized Functional Combinatory Categorial Grammar (GFCCG) – a generalized context free grammar. The authors' conclusion is that this kind of grammar is necessary for some of the level 2 operations, like hierarchical clustering and estimation of situation trees (in their case they capture force objects, actions and groupings). In general, the structure of the nodes in the tree is defined by an ontology. For each situation type a different grammar is needed. However, the generality of the approach comes from the fact that, given the grammar, the control structure of the interpretation algorithm is given by a (general purpose) algorithm for language recognition (parsing). This grammatical approach is very close to the rule based approach. However, this approach is claimed to be superior due to the fact that "the process of rule instantiation in a first order logic reasoning system is similar to instantiating a context-free grammar rule within a parser, but there is a fundamental difference. The resulting object in a parser spans a set of tokens (the leaf nodes of the tree of which it is the root) and cannot be combined with an object that spans any of them". In other words, the advantage of grammatical approach comes from some restrictions on the expressiveness of the syntax of the language (i.e., the grammar).

To support the claims, the paper provides two examples of assessment of the force deployment events. Moreover, the paper provides results of experiments whose objective was to establish the performance of the proposed approach in terms of the accuracy of the interpretation of histories. Additionally, the paper discusses some optimizations that are useful for the purpose of the scaling of the algorithms the authors are working with.

The fourth paper is *Designing Ontologies for Higher Level Fusion*, by Little and Rogova. The basic presumptions in this paper are that "ontologies offer a necessary framework for reasoning about situations" and that "A comprehensive metaphysical approach is required for the construction of ontologies, particularly at the upper-most levels, which can provide consistent and comprehensive models of reality that formally describe the kinds of complex relation-types needed for reasoning about the complex entities found in higher-level fusion applications". The main goal of the paper is then to propose a framework for building formal ontologies that are compatible with the accepted metaphysical assumptions and to provide a process for achieving such a goal.

While a number of upper level ontologies exist, e.g., DOLCE or SUMO, the authors have chosen the Basic Formal Ontology (BFO) founded by Barry Smith (http://www.ifomis.org/bfo). BFO consists of sub-ontologies. This paper uses two of them: SNAP and SPAN. SNAP is used to capture various structural views of the world (continuants) at a time instant. SPAN, on the other hand, is used to represent processes (occurrents).

The main focus of the paper is a proposal of a six step process for constructing BFO-based ontologies, with examples from the project in which this methodology was used (situation assessment in a post-disaster environment). In the context of this project, examples of relations from both SNAP and SPAN, as well as between the two types, are given. As reported by the authors, the project was a large scale exercise in building an ontology for disaster relief through which they collected many definitions of terms classified into SNAP and SPAN.

Situations in this paper are understood as collections of relations relevant to a specific view. Stress is put on the formal classification of relations into mereology (part-whole relations) and topology (the nature of space). These two kinds of relation have mathematical theories behind them, which gives them the property of generality, since they are applicable to various domains.

The paper *Ontology-Based Situation Awareness*, by Kokar et al. first differentiates computer awareness from human awareness. Then the paper focuses on computer awareness. The two main aspects of this paper are "situation" and "awareness". Following ideas of Barwise, situation is treated as a first-class object. This is realized by devising an ontology (Situation Theory Ontology, or STO for short) in which situation is a class, and moreover, various situations can be of different situation types. Consequently, in STO one can represent both specific instances of situations (instances of the class Situation) and whole classes of situation (situation types).

The awareness aspect means that in order to be aware, the agent must be able to infer various facts about a situation – facts that are not explicitly given or represented. The necessity of the inference capability for situation awareness comes from the fact that relations, which are the focus of awareness, are not given by perception, but they often need to be inferred from both the given facts and the background knowledge. Since the stress is put on computer support for human awareness, the subject of this paper is how to implement computer awareness which, as mentioned above, requires the capability of inference in a computer. To achieve this goal, computer processable semantics is needed. This leads to the need for formal languages with computer processable semantics. While human situation awareness needs to be measured and supported, computer awareness needs to be defined and implemented.

The theoretical foundations for this work have been given by Barwise, Perry and Devlin in the form of Situation Theory. The authors of this paper then have captured a significant part of this theory and formalized them in the language of OWL, with some parts expressed in terms of rules. The paper provides an extensive discussion of how the developed ontology matches Situation Theory.

To explain the main implications of the proposed approach, a simple example is discussed, including the representation of some aspects of the example in terms of OWL. In particular, it is shown what kind of inference a computer can perform using STO, how instances of the Situation class can be captured, and how computers can exchange situation objects. The example shows that the same scene can be viewed as many different situations; it all depends on whose point of view is considered. In the scene where a dog, a cat, a mouse and a cat owner are involved, there are at least four views of situation. The owner is concerned about the safety of his cat and thus focuses on the aspects of the scene in which his cat could be hurt. The dog is looking for ways to catch the cat. The cat is focused on the escape routes. The mouse is all happy for as long as the cat is busy.

STO is offered as a unifying ontology for situation awareness. The full version of this ontology is posted on the Web so that anybody can access it. Since it is expressed in a formal language, the semantics is clear and thus the OWL STO representation can serve as an unambiguous reference model for future development of an ontology for situation awareness.

The paper *Factored Reasoning for Monitoring Dynamic Team and Goal Formation*, by Pfeffer et al. discusses an approach to detecting one kind of situation – a threat (an ambush), i.e., a situation in which a loosely coupled team consisting of a number of units (agents) is moving towards a specific target with the intent to attack the target. Thus the main issue is to detect the intents of the units and of the team (the goals). In order to achieve such an objective, the monitoring of a number of variables provides input – positions of the units (tracking), the team structure (which unit belongs to which team) and the communication among the units.

The authors state that for this kind of task the "natural" approach to inference in such a model would be to consider the whole space of the problem at once and use particle filtering (PF). However, the high dimensionality of the problem would cause the probabilities of particles having good position to be very small and thus would lead to high errors. Consequently, the authors propose an approach in which the system combines particle filtering with the strategy of reasoning locally about unit positions and globally about team structure and goals.

The important aspect of the algorithms presented in this paper is that situations (team structures and goals, as well as the set of actions that realize the same global goal) change over time. Although each unit has its own goals, it also realizes the team's goal. Another aspect of the approach is the use of a model for representing constraints and dependencies among the system variables. Constraints represent possible movement directions and the desirability of a particular move. For representing dependencies dynamic Bayesian networks are utilized. The model represents how next states of the units depend on the current goals and previous positions, while satisfying the given constraints. The overall goal of the system is to infer a probability distribution over the variables that represent unit positions, team structure and team goals.

The approach is tested on simulated scenarios of urban warfare. Constraints are provided by a map in which the action takes place. The results show that the approach can track up to 20 units and that their algorithm performs better than standard particle filtering and than an algorithm that performs all reasoning locally. The metrics used in the evaluation of the algorithm were *precision* (the fraction of threats reported correctly by the algorithm as threats) and *recall* (the fraction of real threats detected by the algorithm).

The last paper in this special issue (*High Level Information Fusion for Tracking and Projection of Multistage Stealthy Cyber Attacks*, by Yang et al.) reports on a study of the use of information fusion in the domain of cyber security – intrusion detection and tracking of cyber attacks as well as the projection of the attacks and impact assessment. Referring to the JDL model [1], this paper proposes a correspondence between the intrusion detection domain and the JDL model. In this respect, the paper addresses the problems of level 1 processing (tracking) as well as level 3 (threat projection and impact assessment). The paper also proposes a correspondence with the Endsley's model of situation awareness [2]. Situation assessment is implicit in the approach.

It is assumed that an existing Intrusion Detection System (IDS) provides input (alerts) to the fusion system described in this paper. Additionally, two models are provided to the system – the Potential Attack Sequence Template and the Information Exposure Sequence Map, which "are to represent all potential sequences of attack methods a cyber attacker can use and the orders over which vulnerabilities may be exposed due to network and system configurations, respectively". Thus two aspects are separated in the models: the attacker behavior and the network (network entity exposed). These two reference models capture the expectations of the behaviors and are used as templates for pattern matching. The authors state that the fusion system makes use of both Dempster-Shafer and probabilistic approaches to deal with the uncertainty.

Alerts are processed by the subsystem called INFERD whose goal is to correlate incoming alerts with tracks. It is a challenging task due to the large amounts of data and the necessity to process this information in real time. The second subsystem (TANDI) assesses which network entity is most likely to be compromised next.

INFERD was evaluated on a set of scenarios prepared by an independent source, i.e., not by the research team. The network configuration was also provided. The metrics used in the evaluation of INFERD were precision, recall, fragmentation and mis-association. The paper shows the results of these experiments – the metrics plotted against the classification threshold (the degree of match to the ground truth). Based on these results, the authors conclude that they provided "good "confidence" in identifying cyber attacks".

TANDI was evaluated on simulated data with randomly generated attacks (three sets) on artificially created networks. The reason for not using the same data as for INFERD was that more detail about the network was needed than provided by the external source. The scenarios included attacks that were not covered by the templates. The results of the experiments are shown as aggregation of five metrics related to, among others, false positives, false negatives, abnormality. The conclusion is that the results "exhibit promising performance with TANDIs approach to model separately the attacker behavior and the network vulnerabilities".

The papers selected for publication in this special issue provide coverage for many of the topics listed at the beginning of this editorial. However, first, many of these topics have not been addressed, indicating that either these topics are not attacked by the research community and thus more research in this area is needed, or simply because there were no submissions of this kind to this special issue. Second, some of the topics have been addressed in more than one of the papers. And third, some of the solutions proposed in the papers address the same problem in a different way. While these observations indicate that there is no single, homogeneous theory of high-level information fusion, the diversity of the approaches presented in the papers shows that the community is not locked within a narrow paradigm, but rather is looking to many different ways to solve the problems posed to it by the applications.

References

[1] A.N. Steinberg, C.L. Bowman, F.E. White, Revisions to the JDL data fusion model, in: The Joint NATO/IRIS Conference, 1998.
[2] M.R. Endsley, Toward a theory of situation awareness in dynamic systems, Human Factors Journal 37 (1) (1995) 32–64.

Mieczyslaw M. Kokar
Northeastern University Boston, MA 02115, United States
Tel.: +1 617 373 4849; fax: +1 617 373 8970
E-mail address: mkokar@ece.neu.edu

Gee Wah Ng
DSO National Laboratories, Singapore 118230, Singapore
E-mail address: geewah.ng@gmail.com