

This document is published in:

Bossé, E. et al. (eds.), 2013. *Prediction and Recognition of Piracy Efforts Using Collaborative Human-Centric Information Systems*, Proceedings of the NATO Advanced Study Institute (ASI), Salamanca, Spain 19-30 September, 2011, (NATO Science for Peace and Security, Sub-Series - E: Human and Societal Dynamics, v. 109), IOS Press, pp. 176-183.

DOI: 10.3233/978-1-61499-201-1-176

© 2013 IOS Press

A Practical Approach to the Development of Ontology-Based Information Fusion Systems

Juan Gómez-Romero¹, José M. Molina, Miguel A. Patricio, Miguel A. Serrano
Applied Artificial Intelligence Group, University Carlos III of Madrid, Spain

Abstract. Ontology-based representations are gaining momentum among other alternatives to implement the knowledge model of high-level fusion applications. In this paper, we provide an introduction to the theoretical foundations of ontology-based knowledge representation and reasoning, with a particular focus on the issues that appear in maritime security –where heterogeneous regulations, information sources, users, and systems are involved. We also present some current approaches and existing technologies for high-level fusion based on ontological representations. Unfortunately, current tools for the practical implementation of ontology-based systems are not fully standardized, or even prepared to work together in medium-scale systems. Accordingly, we discuss different alternatives to face problems such as spatial and temporal knowledge representation or uncertainty management. To illustrate the conclusions drawn from this research, an ontology-based semantic tracking system is briefly presented. Results and latent capabilities of this framework are shown at the end of the paper, where we also envision future opportunities for this kind of applications.

Keywords. Information Fusion, Knowledge-Based Systems, Ontologies, Maritime Security

Introduction

Data and information fusion (DIF) research area studies theories and methods to effectively “combine data from multiple sensors and related information to achieve more specific inferences that could be achieved by using a single, independent sensor” [1]. The widely-accepted Joint Directors of Laboratories (JDL) model classifies fusion processes into five operational levels corresponding to different stages of the transformation from input signals to decision-ready knowledge [2, 3]; namely: signal feature assessment (L0), entity assessment (L1), situation assessment (L2), impact assessment (L3), and process assessment (L4).

Low-level data fusion, corresponding to JDL L0 and L1 levels, has received considerable attention during the last decades, which has resulted in a myriad of theories, algorithms and tools to process multi-sensor signals and to estimate object properties. These approaches have been successfully applied to several domains, such as radar-based tracking, video surveillance, and ambient intelligence. On the other hand, high-level fusion procedures, corresponding to JDL L2 and L3, aim at obtaining a description of the relations between the objects in the scenario. These relations are expressed in symbolic terms (actions, intentions, threats), instead of the numerical measures (density functions, movement vectors) computed in L0 and L1. The ultimate objective of

¹ Corresponding author: Juan Gómez Romero, Applied Artificial Intelligence Group, University Carlos III of Madrid. Avd. de la Universidad Carlos III 22, 28270 Colmenarejo, Madrid, Spain; E-mail: jgromero@inf.uc3m.es

high-level fusion procedures is to achieve situation assessment; i.e., to *understand* the scene in terms that can be easily communicated to the intelligence officer, to *evaluate* short and long-term threats, and to *support* decision making.

Unfortunately, high-level fusion is a problem still far from being solved. High-level fusion requires systems to process and to interpret abstract information, thus exhibiting abilities close to human cognition. In addition, modern fusion applications must be able to work in problems where the world-behavior is very complex and unpredictable, and where contextual influences are important or even critical. This requires the implementation of flexible and dynamic situation models, able to adapt to unexpected situations, as well as the exploitation of context knowledge to incorporate contextual effects to the systems.

For these reasons, symbolic formalisms have been proposed to represent and reason with high-level information. Cognitive approaches propose building a symbolic model of the world, expressed in a logic-based language, to abstractly represent scene objects and their relations. Cognitive approaches are more robust and extensible than quantitative proposals, but they require the development of suitable interpretation and reasoning procedures, which is not assumable or even possible in all cases. In addition, cognitive models must implement procedures to bridge the gap between abstract representations in the symbolic language and concrete measures acquired by sensors, which is known as the grounding problem.

Ontologies have recently received a considerable attention as proper formalisms to create symbolic models in high-level fusion systems [4]. An ontology, in the knowledge engineering area, is defined as “an explicit formal specification of how to represent the objects, concepts and other entities that are assumed to exist in some area of interest and the relationships that hold among them” [5]. Ontologies support formal information representation and reasoning while promoting knowledge reuse. These properties make ontologies very suitable in high-level fusion, which entails the use of a common communication language between the actors involved in the process, and the integration of several heterogeneous information sources. Unfortunately, current procedures and tools for the development of ontology-based fusion systems are not fully standardized, or even prepared to work together in medium-scale system.

In previous research works, we have presented an ontology-based framework for contextual interpretation of data acquired from a visual sensor network [6]. The framework constructs a symbolic model of the scene by integrating tracking data and contextual information. The scene model, represented with formal ontologies, supports the execution of reasoning procedures in order to: (i) obtain a high-level interpretation of the scenario; (ii) provide feedback to the low-level tracking procedure to improve its accuracy and performance. In the current paper, we discuss some lessons learned during this development, and we study the applicability of our conclusions to the harbor scenario.

The remainder of this paper is structured as follows. Section 1 provides a brief introduction to the notion of ontology and the benefits contributed by ontologies to high-level fusion systems, and describes some ontology features that can be exploited in the maritime domain, such as knowledge representation and exchange, entity classification, management of spatial knowledge, and rule-based reasoning. Section 3 depicts the architecture of a general ontology-based high-level fusion system, as well as some specific details of the framework presented in [6]. Section 4 discusses some details of the implementation of the architecture. Finally, Section 5 summarizes the conclusions of the paper and presents some prospective directions for future work.

1. Knowledge Representation and Reasoning with Ontologies

1.1. Principles

An ontology is a knowledge model which describes the objects in an application domain from a common perspective by using a language that can be automatically processed. This language is usually a Description Logic-based representation [7].

The use of ontologies in DIF results in several advantages: (i) abstract representation of information, which improves interpretability of the system and make it easier for the user to interact with it; (ii) reasoning with logic-based formalisms, which allows inferring new knowledge; (iii) extensibility of the knowledge bases, which facilitates the application of the model in diverse application domains; (iv) standardization, which supports interoperation between different modules and systems.

The basic ontological representation primitives are concepts, relations, instances and axioms. Concepts or classes (noted with capital letters C , D) represent the basic ideas of the domain that must be apprehended, and they determine sets which classify domain objects. Instances or individuals (noted a , b) are concrete occurrences of a concept. Relations or roles (r , s) represent binary connections between individuals or individuals and typed values (integers, strings, etc.). Axioms (τ , φ) establish restrictions over concepts, instances, and relations, describing their attributes by delimiting their possible interpretation. Axioms involve atomic or complex concepts and relations, which can be composed by recursively applying the constructors allowed by the logic. DLs are named after the list of allowed constructors with a string of capital letters. For example, OWL 2 –the standard Ontology Web Language– is almost equivalent to the $\mathcal{SRQ}(D)$ logic [8]. Ontological descriptions are usually created with a proper ontology editor. Protégé² (open-source, free) and TopBraid Composer³ (commercial) are two of the most recognized ontology development tools.

Reasoning with ontologies is an automatic procedure that infers new axioms that have not been explicitly included in the knowledge base but are logical consequences of the represented axioms. Generally speaking, an axiom τ is entailed by an ontology K (noted $K \models \tau$) if every possible realization of K satisfies τ . The basic reasoning task regarding ontology concepts is concept satisfiability. Intuitively, a concept is satisfiable if it is not contradictory of the rest of the knowledge in the ontology. Another important task is concept subsumption, which infers if a concept D is more general than another concept C ($C \sqsubseteq D$). Similarly, the basic inference task with ontology individuals is to test if an axiom is consistent; i.e., the axiom is not contradictory of the other axioms in the ontology, or in particular of instance axioms. If the assert is a membership axiom $a : C$ (meaning a belongs to C), this test is called instance checking. The computational complexity of the reasoning procedures directly depends on the expressivity of the DL language considered –the more expressive is the language, the higher is its complexity.

Reasoning tasks can be transparently executed with DL inference engines (also named reasoners), which allow loading and querying ontologies expressed in the OWL 2 standard language. Pellet⁴ and RACER⁵ are two freely available DL reasoners.

² <http://protege.stanford.edu>

³ http://www.topquadrant.com/products/TB_Composer.html

⁴ <http://clarkparsia.com/pellet/>

⁵ <http://www.racer-systems.com/>

1.2. Ontologies for High-Level Fusion in the Maritime Domain

Situation and threat assessment in the harbor surveillance scenario is based on matching expected with observed and inferred track, object and situational items properties. Several authors have pointed out that context plays a central role in such assessment process. Firstly, context provides additional restrictions to fusion processes, which are in turn used to check the consistency of a situational hypothesis built on observed data. In addition, context can be applied to enrich the situational hypothesis by linking extended information available from own or external knowledge bases.

Ontologies are being successfully applied to create a uniform, widely accessible knowledge model to support contextual high-level fusion in the harbor domain [9, 10, 11]. Ontologies represent harbor zones and vessel classes to describe non-threatening behaviors consistent with normality schemas extracted from IMO or acquired from experienced officers. Standard ontology reasoning procedures can be applied to classify normal behaviors as friendly, leaving unclassified behaviors as suspicious.

Situation assessment procedures must shortly afterwards determine the threat level of these unidentified suspicious behaviors, thus requiring the DIF system to find suitable hypotheses to explain the current inputs. This kind of reasoning is not directly supported by ontologies, but it can be implemented by relying on extended services offered by reasoning engines. Uncertainty can be incorporated at this layer through different paradigms; e.g., the belief-based argumentation system (BAS) [12] –a non-monotonic approach for reasoning under uncertainty that combines symbolic logic with belief theory–, or Bayesian-based *ad hoc* formalisms [13].

Besides representation and reasoning features, the use of ontologies can be very helpful to facilitate knowledge exchange between different entities –in some cases, it has been reported that coastal defense can mobilize more than three different authorities, including maritime port authorities, coast guards and navy. Ontologies support the definition of a common knowledge exchange language independent from the internal procedures of the involved agents. The creation of a central knowledge repository with agreed semantics, as described in Section 3, would support the development of distributed applications for accessing, digesting and presenting this information from a unified perspective.

2. Ontology-Based High-Level Fusion Architecture

Our view on the architectural organization that supports ontology-based high-level fusion is depicted in Figure 1. This schema proposes the implementation of a processing layer on top of the low-level fusion procedures. The context information exploitation layer checks the consistency of current hypothesis with new relevant information, proposes new hypothesis from current low-level data, adds additional facts to make more information accessible to automatic reasoning processes and decision makers, and infers recommendations to improve the performance of the fusion procedures.

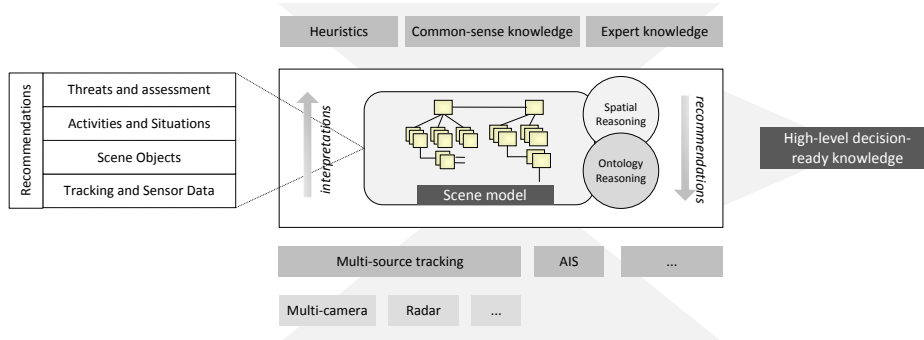


Figure 1. Architecture of Ontology-Based Information Fusion Systems

The architecture conceptually separates low and high-level fusion procedures. On the one hand, low-level fusion mainly concerns multi-source tracking procedures. On the other hand, high-level fusion involves various steps to convert numeric sensor data to symbolic information, as well as other procedures aimed to generate recommendations to adjust the behavior of system components. Fusion is thus regarded as a model-building procedure, which results in the construction of an ontological instantiation that abstractly represents the fused scene. The figure shows the layered structure of the ontological knowledge base that supports the contextual layer, including sub-models to symbolically represent raw sensor data, objects, situations, assessments and recommendations. An important component of the architecture is the spatial reasoning module, which specifically focus on detecting and updating qualitative topological relations of the model. This module uses an auxiliary data structure to optimize the calculations involving geometric entities.

3. Implementation: System Prototype and Tools

In this section, we describe the in-progress implementation of various modules of the previous architecture. The resulting system prototype is being currently tested in video-surveillance [6] and Ambient Intelligence [14] applications. We focus on three main components: the low-level tracker that provides input data to the high-level fusion system, the ontology-based layer managing the scene model and context information, and the spatial reasoning unit.

3.1. Fuzzy Tracking Module

The tracking sub-system used in our prototype implements a video chain with different modules that run in sequence the successive phases of the tracking process on input data provided by a single camera [15]. Accordingly, the input data is the current frame of the video stream, whereas the output data is tracks position and size. Interestingly enough, the tracking module could be replaced by other module able to provide basic position data without any significant modification of the architecture.

The current implementation allows user to select the algorithm that must be applied at each stage of the tracking process. In particular, the tracking system includes an

association procedure named Fuzzy Region Assignment (FRA) [16]. The algorithm integrates visual information at several levels of granularity: low-level image segmentation operations, medium-level smoothness criteria on target features, and high-level constraints on tracking continuity. FRA is based on a Bayesian formulation to determine when a blob is related with a track through an estimated probability. Four heuristic functions (overlap, deform, density, conflict) are used to update the track situation and its dimensions. A set of fuzzy rules derived from experimentation infers the resulting confidence output to define the final association.

3.2. *Ontology-Based Representation and Reasoning*

The scene model includes a hierarchy of ontologies organized according to the JDL abstraction levels: tracking entities, scene objects, activities, and assessments. The ontological model merges *a priori* scene data given by users with sensor data coming from different inputs: video-based tracking, AIS, radar, etc.

Incoming data is managed by using the OWL API [17], a Java programming interface to deal with ontologies, whereas the RACER inference engine is used to load and reason with the scene model. RACER has been chosen because it includes support for different kind of inference rules such as deductive, abductive, spatial, temporal, etc. In particular, abductive rules in the nRQL (new RACER query) language are defined to create higher-level information from lower-level data. These rules make intensive use of spatial knowledge represented with the Region Connection Calculus formalism (RCC), which is also supported by RACER in the form of a substrate.

The system can manage current and past scene information. This implies that a temporal dimension can be applied in specific rules; for example, it is possible to create rules that trigger if “a vessel is stopped in a restricted area during the last ten or twenty time intervals”. However massive storage of historical data may significantly degrade system performance, since the inference engine has to search through a larger number of axioms. Thus, a compromise between data storage and query performance must be implemented; for example, by restricting the temporal window allowed for past events.

3.3. *Spatial Data Management – Dynamic RCC*

Scalability of the model can be seriously compromised when several objects are in the scene at the same. To avoid this problem, a specific module for spatial data management has been implemented, namely Dynamic RCC [19]. Dynamic RCC has two objectives: (i) representation and reasoning with spatial properties in the ontological model; (ii) efficient instantiation and update of spatial properties of detected objects.

Dynamic RCC encompasses three main components: a knowledge base with the spatial properties of individuals corresponding to scene objects; an optimized geometric model, encompassing a geometric model and an auxiliary data structure; and a RCC implementation that stores qualitative spatial relationships. Dynamic RCC proceeds as follows. First, geometric features of static and dynamic objects, represented at object level, are processed and inserted into the geometric model. Next, a full topological analysis between the new/updated and the previous geometries is performed. The new topological relations which change from the previous state are then updated in the RCC implementation and sent back to the scene model –actually, changes in the ontologies of the scene model are not necessary, because instance properties and topological relations are stored in a separate substrate.

The geometric model represents spatial 2-dimensional entities in a Euclidean plane. This model is implemented with the Java Topology Suite⁶ (JTS) according to the OpenGIS Simple Features standard⁷, a standard for digital storage of geographical data. OpenGIS defines a set of methods to evaluate the spatial relationships, a set of methods to support spatial analysis, relational operators between entities, and several kinds of representations. Although OpenGIS spatial predicates and RCC are not directly compatible, they can be easily mapped.

The topological analysis has a quadratic complexity, since it is necessary to make a pairwise comparison between all the geometric entities of the scene. Therefore, it is convenient to reduce the comparisons only to those geometries that are candidates to modify the spatial relations of an object. The spatial data structure maintains a hierarchical topological sort on the Euclidean space of the scene objects to support the retrieval of the possible candidate geometries involved in a topological analysis.

4. Discussion and Future Work

In this work, we have studied the contributions and the advantages of ontologies as the knowledge representation formalism of high-level information fusion procedures. Ontologies support the creation of a symbolic scene model that serves as a common repository to be exploited by different applications. Reasoning capabilities of ontologies can be used to implement high-level fusion procedures aimed at interpreting the current situation and automatically determining existing threats to help decision makers. We have also shown some specific problems that must be solved in the practical implementations of an ontology-based framework for DIF.

The system presented in this paper has been applied to solve high-level fusion problems in video-based applications. Lessons learnt in that domain can be applied in the harbor scenario. In particular, encoding specific rules and restrictions, as well as incorporating relevant context knowledge, would require the participation of experts. Nevertheless, the abstract representation features of ontologies would simplify this process. In addition, ontologies facilitate knowledge integration and reuse, which would be very useful to reduce the effort required to adapt the system to different port configurations and to incorporate multi-source or multi-modal information.

More discussions and implementations of ontology-based high-level fusion systems will be certainly useful to foster the creation of competitive frameworks. The first step towards a reliable framework for real use is to develop a practical implementation focused on the specifications of a concrete harbor. We strongly believe that the architecture and the technologies presented in this paper are the blueprints and the tools that will support these future developments.

Acknowledgements

This research activity is supported in part by Projects CICYT TIN2008-06742-C02-02/TSI, CICYT TEC2008-06732-C02-02/TEC, CAM CONTEXTS (S2009/TIC-1485) and DPS 2008-07029-C02-02.

⁶ <http://www.vividsolutions.com/jts/JTSHome.htm>

⁷ <http://www.opengeospatial.org/standards/sfs>

References

- [1] Llinas J, Hall DL (2009) Multisensor data fusion. In: Handbook of multisensor data fusion (eds. Liggins ME, Hall DL, Llinas J). CRC Press, pp 1–14
- [2] Steinberg AN, Bowman CL (2004) Rethinking the JDL data fusion levels. In: Proceedings of the MSS national symposium on sensor and data fusion. Columbia, SC, USA
- [3] Llinas J, Bowman CL, Rogova G, Steinberg AN, Waltz E, White F (2004) Revisiting the JDL data fusion model II. In: Proceedings of the 7th international conference on Information Fusion (Fusion 2004). Stockholm, Sweden, pp 1218–1230
- [4] Nowak C (2003) On ontologies for high-level information fusion. In: Proceedings of the 6th Int. Conference on Information Fusion (Fusion 2003). Cairns, Australia, pp 657–664
- [5] Staab S, Studer R (2009) Handbook on Ontologies. Springer-Verlag
- [6] Gómez-Romero J, Patricio MA, García J, Molina JM (2011) Ontology-based context representation and reasoning for object tracking and scene interpretation in video. *Expert Systems with Applications*, 38(6), 7494–7510
- [7] Baader F, Horrocks I, Sattler U (2008) Description Logics. In: Handbook of knowledge representation (eds. van Harmelen F, Lifschitz V, Porter B). Elsevier, pp. 135–180
- [8] Hitzler P et al. (2008) OWL 2 Web Ontology Language primer. <http://www.w3.org/TR/owl2-primer/> (accessed Jul. 2011)
- [9] García J, Gómez-Romero J, Patricio MA, Molina JM, Rogova G (2011) Representation and exploitation of context knowledge in a harbor surveillance scenario. In: Proceedings of the 14th Int. Conference on Information Fusion (Fusion 2011). Chicago, IL, USA
- [10] García J, Gómez-Romero J, Rogova G (2011) Contextual knowledge and information fusion for maritime piracy surveillance. In: Proceedings of the NATO-ASI Prediction and Recognition of Piracy Efforts using Collaborative Human-Centric Information Systems. Salamanca, Spain
- [11] van den Broek AC, Neef RM, Hanckmann P, van Gosliga SP, van Halsema D (2011). Improving maritime situational awareness by fusing sensor information and intelligence. In: Proceedings of the 14th Int. Conference on Information Fusion (Fusion 2011). Chicago, IL, USA
- [12] Rogova G, Scott P, Lollett C, Mudiyanur R (2006) Reasoning about situations in the early post-disaster response environment. In: Proceedings of the 9th Int. Conference on Information Fusion (Fusion 2006). Florence, Italy
- [13] Carvalho RN, Haberman R, Costa PCG, Laskey KB, Chang KC (2011). Modeling a probabilistic ontology for maritime domain awareness. In: Proceedings of the 14th Int. Conference on Information Fusion (Fusion 2011). Chicago, IL, USA
- [14] Gómez-Romero J, Serrano MA, Patricio MA, García J, Molina JM (In Press) Context-based scene recognition from visual data in smart homes: an Information Fusion approach. *Journal of Personal and Ubiquitous Computing*
- [15] Besada JA, García J, Portillo J, Molina JM (2005) Airport surface surveillance based on video images. *IEEE Transactions on Aerospace and Electronic Systems* 41, 1075–1082
- [16] García J, Patricio MA, Berlanga A, Molina JM (2009) Fuzzy region assignment for visual tracking. *Soft Computing*, 15(9), 1–20
- [17] Horridge H, Bechhofer S (2009) The OWL API: A Java API for working with OWL 2 ontologies. In: Proceedings of the 6th OWL Experienced and Directions Workshop (OWLED 2009). Chantilly, VA
- [18] Serrano MA, Patricio MA, García J, Molina JM (2011) A general purpose context reasoning environment to deal with tracking problems: an ontology-based prototype. In: Proceedings of the 6th Int. Conference on Hybrid Artificial Intelligent Systems (HAIS 2011). Wroclaw, Poland, pp. 144–154
- [19] Serrano MA, Gómez-Romero J, Patricio MA, García J, Molina JM (2011) Topological properties in ontology-based applications. In: Proceedings of the 11th Int. Conference on Intelligent Systems and Applications (ISDA 2011). Córdoba, Spain