

This document is published in:

2010 13th Conference on Information Fusion (FUSION 2010):

Edinburgh, UK. 26-29 July 2010.

© 2010 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

Strategies and Techniques for Use and Exploitation of Contextual Information in High-Level Fusion Architectures

J. Gómez-Romero^a, J. García^a, M. Kandefer^c, J. Llinas^{b,c},
J.M. Molina^a, M.A. Patricio^a, M. Prentice^c, S.C. Shapiro^c

jgromero@inf.uc3m.es, jgherrer@inf.uc3m.es, mwk3@buffalo.edu, llinas@buffalo.edu,
molina@ia.uc3m.es, mpatrici@inf.uc3m.es, mjpp44@buffalo.edu, shapiro@buffalo.edu

^a Applied Artificial
Intelligence Group,
University Carlos III of Madrid

^b Chair of Excellence
University Carlos III of Madrid

^c Department of Computer Science
& Engineering
University at Buffalo

Abstract - Contextual Information is proving to be not only an additional exploitable information source for improving entity and situational estimates in certain Information Fusion systems, but can also be the entire focus of estimation for such systems as those directed to Ambient Intelligence (AI) and Context-Aware(CA) applications. This paper will discuss the role(s) of Contextual Information (CI) in a wide variety of IF applications to include AI, CA, Defense, and Cybersecurity among possible others, the issues involved in designing strategies and techniques for CI use and exploitation, provide some exemplars of evolving CI use/exploitation designs on our current projects, and describe some general frameworks that are evolving in various application domains where CI is proving critical.

Keywords: high-level fusion, context, ontologies.

1 Introduction

It can be argued that there are four categories of information that can be applied to any Information Fusion (IF) problem: observational data, a priori knowledge models, inductively learned knowledge, and contextual information. For a broad class of applications, many IF processes and systems have been designed to work largely on the first two types of information; these are the class of systems built on a deductive-model foundation and that largely employ observational data in a scheme to more or less match the data against the models. These approaches can work well for what could be called well-behaved and well-studied problem domains but cannot be expected to work in problems where the “world-behavior” is very complex and unpredictable or in problems where contextual influences are important or even critical. One such class of defense-type applications are those involving Counterinsurgency (“COIN”) and irregular warfare applications, where both complex and unpredictable adversarial behavior can be expected, and where even political and religious contextual effects can be prime

drivers or constraints to such behavior. There have been some studies in the IF community that have attempted to incorporate contextual effects, such as the effect of local terrain on object kinematics in fusion-based tracking system development (e.g., [1]).

Further, in certain applications, the contextual setting is in fact the prime focus of the IF system to be developed, i.e., fusion of observational data and use of knowledge models is employed in order to estimate the contextual framework itself; these are such applications as Ambient Intelligence Systems and Context-aware systems. These applications often involve visual, imagery-based observational data. A related application area thus is that of video-based surveillance systems, where contextual effects are also important.

Teams from the University of Carlos III of Madrid in Spain (UC3M) and the University at Buffalo (UB) in the USA are sharing their interests in this research area and have joined in this paper to describe their respective approaches to representing, using, and exploiting contextual information in these different application domains. This paper therefore has two main sections framed around the research work in the respective domains, with UC3M leading the discussion on ambient intelligence, context-aware, and video surveillance applications, and UB leading the discussion on COIN-type defense applications.

This paper is organized as follows. Section 2 discusses the role of contextual information in high-level IF. We describe two architectural frameworks for context knowledge exploitation, namely the “a priori” and the “a posteriori” models. In Section 3, we present a context-based approach to IF in video surveillance developed at UC3M. This proposal reveals the advantages of ontology-based context representation and reasoning in Ambient Intelligence and Context-aware systems. Section 4 presents an approach to context-based Information Retrieval for COIN developed at UB. Process flow and advantages of the approach are considered. The paper ends with some conclusions and prospective research directions.

2 The role of contextual information in high-level Information Fusion

2.1 Understanding and using context

Contextual Information is that information that can be said to “surround” a situation of interest in the world. It is information that aids in understanding the (estimated) situation and also aids in reacting to the situation, if a reaction is required. It can be seen as a set of constraints to a reasoning process about a situation; Kandefor and Shapiro define it this way [2]: “the structured set of variable, external constraints to some (natural or artificial) cognitive process that influences the behavior of that process in the agent(s) under consideration.” There are of course other definitions of this somewhat ambiguous term, such as that offered by Dey and Abowd, who state that context is “any information (either implicit or explicit) that can be used to characterize the situation of an entity” [3]. Contextual information can be relatively or fully static or can be dynamic, possibly changing along the same timeline as the situation. It is also likely that the full characterization and specification of Contextual Information may not be able to be known at system/algorithm design time, except in very closed worlds.

Thus, we envision an “a priori” framework of exploitation of Contextual Information that attempts to account for the effects on situational estimation of that Contextual Information (CI henceforth) that is *known at design time*; there is a question of the ease or difficulty

involved in integrating CI effects into a fusion system design or into any algorithm designs. This issue is influenced in part by the nature of the CI and the manner of its native representation, e.g., as numeric or symbolic, and the nature of the corresponding algorithm. Strategies for a priori exploitation of CI may thus require the invention of new hybrid methods that incorporate whatever information an algorithm normally employs in estimation (usually observational data) with an adjunct CI exploitation process. Note too that CI may, like observational data, have errors and inconsistencies itself, and accommodation of such errors is a consideration for hybrid algorithm design. In this case then, we have a notional processing operation as shown in Fig 1.

Similarly, we envision the need for an “a posteriori” CI exploitation process, due to at least two factors:

- 1) That all relevant CI may not be able to be known at system/algorithm design time, and may have to be searched for and discovered at runtime, as a function of the current situation estimate.
- 2) That such CI may not be of a type that was integrated into the system/algorithm designs at design time and so may not be able to be easily integrated into the situation estimation process.

In this case then we envision that at least part of the job of posteriori CI exploitation would be of a type that checks the consistency of a current situational hypothesis with the newly-discovered (and situationally-relevant) CI, but also – if the hypothesis is in fact consistent – adds some explanatory aspects to the declared hypothesis.




Figure 1. Notional Process Flow for “A Priori” CI Exploitation

There are yet other system engineering issues. The first is the question of accessibility; CI must of course be accessible in order to use it, but accessibility may not be a straightforward matter in all cases. One question is whether the most-current CI is available; another may be that some CI is controlled or secure and may have limited availability. The other question is one of representational form. CI data can be expected to be of a type that has been created by “native” users – for example weather data, important in many fusion applications as CI, is generated by meteorologists, for meteorologists (not for fusion system designers) – thus, even if these data are available, there is likely to be a need for a “middleware” layer that incorporates some logic and algorithms to both sample these data and shape them into a form suitable for use in fusion processes of various type. In even simpler cases, this middleware may be required to reformat the data from some native form to a useable form.

In spite of some a priori mapping of how CI influences or constrains the way in which situational inferences or estimates can be developed, which may serve certain environments, the defense and security type applications, with their various dynamic and uncertain types of CI, demand a more adaptive approach. Given a nominated situational hypothesis H_f from a fusion process or “engine” (that might already have accounted for some CI, as explained above), the first question is: what CI type information is relevant to this hypothesis? As cited by Kandefer and Shapiro in [1], “The relevancy problem is defined by Ekbia and Maguitman as “the problem of identifying and using properly [only] the information that should exert an influence on our beliefs, goals, or plans” [4]. Said otherwise, relevant CI is only that information that influences our interpretation or understanding of H_f . Presuming a “relevancy filter” can be crafted, a search function would explore the available or retrievable CI and make this CI available to a “posteriori” reasoning engine. That reasoning engine would then use a CI-guided subset

of Domain Knowledge, and the retrieved CI to reason over H_f to first determine consistency of H_f with the relevant CI. If it is inconsistent, then some type of adjudication logic will need to be applied to reconcile an inconsistency between the fusion process that produced H_f and the posteriori reasoning process that judges it as inconsistent. If however H_f is judged as consistent with the additional CI, an expanded interpretation of H_f could be developed, providing a deeper situational understanding. This processing flow is depicted in Fig 2.

As noted in Fig 2, this overall process, which can be considered a “Process Refinement” operation, would be a so-called “Level 4” process in the context of the JDL Data Fusion Process Model [5], which identifies Level 4 as the class of adaptive operations targeted to fusion process enhancement.

2.2 Related works in CI exploitation in Information Fusion

Symbolic models have been applied to acquire, represent, and exploit knowledge in fusion, and particularly in fusion of visual information. In the last years, the interest in ontologies has increased considerably [6], and its use is becoming more and more frequent. Nevertheless, despite of the fact that most of the current approaches combine contextual and perceptual information, they do not explicitly describe how context is characterized and integrated in the fusion process. This paper is aimed to encourage the discussion on these topics.

Previous ontology-based fusion researches can be classified according to the four levels defined by the JDL model. At image-data level (i.e., JDL level 0), one of the most important contributions is COMM (Core Ontology for MultiMedia), an OWL ontology to encode MPEG-7 data [7]. Similarly, the Media Annotations Working Group of the W3C is working in an OWL-based language for adding metadata to Web images and videos [8].




Figure 2. Notional Process Flow for “A Posteriori” CI Exploitation

Other proposals are targeted at modeling video content at object level (i.e., JDL L1). For example, a framework for video event representation and annotation is described in [9]. This framework includes two languages, namely VERL (Video Event Representation Language) and VEML (Video Event Markup Language), used to describe processes and to markup video sequences, respectively. Other authors have discussed and improved this approach to support the representation of uncertain knowledge [10]. The research work in [11] presents a symbolic representation for the data managed by a tracking algorithm. Similarly, recently it has been a first approximation to the development of an ontology for tracking data [12].

Scene interpretation issues (i.e., JDL L2 and L3) are being dealt with ontologies as well. In [13], it is presented a proposal for scene interpretation based on Description Logics and supported by the reasoning features of RACER inference engine. The problem of representing high-level semantics of situations with a computable formalism is also faced in [14], where an OWL ontology encoding Barwise's situation semantics is developed.

All these research works focus on contextual scene recognition, but as previously mentioned, it is also interesting to apply this knowledge to refine image-processing algorithms (which corresponds to JDL L4), and particularly trackers. A preliminary approach to this topic has been presented in [19]. In the present paper, some aspects of the current state of this research are reviewed, and its contributions from an architectural and knowledge management point of view are discussed.

3 Applications in Video Surveillance, Ambient Intelligence, and Context-aware Systems

Third-generation surveillance applications are considered to be the next step in multi-camera security systems [15, 16]. This term designate systems that resemble the nature of the human intelligent process of surveillance, which activates certain cognitive abilities, and that satisfy the requirements of modern surveillance, which are among others the management of a large number of cameras, the geographical spread of resources and the need of many monitoring points. The ultimate goal of third-generation surveillance systems is to automatically achieve a high degree of understanding of the scene from multiple observations to barely require operator attention while cutting component cost.

The existence of multiple cameras requires further researches to develop information fusion procedures to integrate data generated at different locations and reasoning techniques to obtain a high-level and global

interpretation of the scene. A solution to overcome these issues is to provide the image-processing algorithms with additional information about the observed entities not directly obtained by the cameras; i.e., to incorporate CI in the data and information fusion process.

In this domain application, we consider that any external piece of knowledge used to complete the quantitative data about the scene computed by straightforward image-analysis algorithms can be considered context. This definition is coherent with the definitions provided above. Four sources of context knowledge must be taken into account [17]: (i) the scene environment: structures, static objects, illumination and other behavioral characteristics, etc.; (ii) the parameters of the recording: camera, image, and location features; (iii) historic information: past detected events; (iv) soft information provided by humans.

Our approach to the notion of context is intentionally imprecise. Usually, it is difficult to clearly distinguish which information is context and which information is perception, either from a computational or a cognitive perspective. Hence, we propose to create a common symbolic representation and reasoning model to support context exploitation. CI is applied to support two abductive reasoning tasks in our Computer Vision framework: high-level interpretation of scenes and feedback generation to improve performance of low-level image processing procedures.

3.1 An architecture for high-level CI-based visual Information Fusion

The architecture of our Computer Vision framework is depicted in Figure 3. The schema shows the tracking system (the GTL, general tracking layer) and, built upon it, the context-based extension (the CL, context layer). The GTL and the CL communicate through the GTL/CL interface.

The GTL is a classical tracking procedure. The CL manages the CL model, i.e. the ontological representation of the scene (including context and perceived knowledge), and carries out procedures to update and reason with it. The CL model is implemented as a set of interrelated ontologies organized according to the JDL model. The terminological axioms of the CL ontologies establish a controlled vocabulary to describe scenes. The current state of the scene is represented with instances of these concepts and relations. The GTL/CL interface, which guarantees interoperability and independence between both layers, includes methods to update the CL model according to the GTL tracking information, and to consult recommendations calculated by the CL.




Figure 3. Architecture of the Context-Based Computer Vision framework

3.2 An ontological representation of CI

The knowledge of the CL model is structured in four layers, as depicted in Figure 4:

- Camera data (L0). Sequence provided by the cameras (in some processable video format).
- Tracking data (L1). Output of the tracking algorithm represented with ontological terms: frames, tracks and track properties (color, position, velocity, etc.).
- Scene objects (L1-L1/2). Objects resulting from making a correspondence between existing tracks and possible scene objects. For example, a track can be inferred to correspond to a person (possibly by applying context). Scene objects also include static elements, which may be defined a priori.
- Activities (L2). Description of relations between objects which last in time, for example, grouping, approaching, or picking/leaving an object.
- Impacts and threats (L3). Cost or threat value assigned to activities.
- Feedback and process improvement (L4). Abstract representation of the suggestions provided to the tracking procedure.

Each one of these levels corresponds to an ontology in the CL knowledge representation model, namely: TREN, SCOB, ACTV, IMPC and RECO. We provide a reference version of these ontologies¹ in the ontology language OWL that must be refined in a concrete application. For instance, in a particular application it may be interesting to define a concept to represent a Mirror as a specialization of StaticObject and ReflectingObject, which are concepts

of SCOB. For a more extensive description of these ontologies, see [18].




Figure 4. Multi-level Representation of Visual Knowledge

¹ <http://www.giaa.inf.uc3m.es/miembros/jgomez/ontologies/>

The ontologies of the CL model only contain descriptive knowledge; i.e., axioms defining concepts and relations. They offer a vocabulary to describe scenes that must be instantiated in a concrete application according to the entities appearing in the scene, both static and dynamic.

Before starting the analysis of a video sequence, it is necessary to annotate the static elements of the scenario. Annotating the scenario means to create instances of the ontologies describing its properties: object position and size, possible occlusions, enter and exit zones, or any other contextual knowledge that we may consider necessary (see Figure 5.(a)). Additionally, contextual rules must be loaded into the reasoner (see next section). Ontology instances corresponding to the dynamic entities of the scene are created and deleted during the processing of the sequence as a result of the track information provided by the GTL and additional reasoning processes (see Figure 5.(b)).

3.3 Abductive reasoning for scene understanding and feedback

Standard ontology reasoning procedures can be performed within the CL ontologies to infer additional knowledge from the explicitly asserted facts (tracking data and a priori context). By using an inference engine, tasks such as classification or instance checking can be performed. Nevertheless, monotonicity of ontology languages forbids adding new knowledge to the models while reasoning, which is required in scene interpretation.

Scene interpretation is a paradigmatic case of abductive reasoning, in contrast to the Description Logics (DLs) deductive reasoning: abductive reasoning takes a set of facts as input and finds a suitable hypothesis that explains them (sometimes with an associated degree of confidence or probability). This is what is needed in our case: we want to figure out what is happening in the scene from the observed and the contextual facts. In terms of the architecture of the CL, scene interpretation can be seen as an abductive transformation from instances of a lower level ontology to instances of a higher level ontology. Abductive reasoning is not supported by ontologies [19], but it can be simulated by using customized procedures or,

preferably, by defining transformation rules in a suitable query language. RACER², an inference engine for DL ontologies, allows abductive reasoning, and therefore it has been our choice to implement the framework. In this first approximation, we have not considered uncertain abduction, which is a promising direction of future research.

In our framework, abductive rules formally represent contextual, heuristic and common sense knowledge to accomplish high-level scene interpretation and low-level tracking refinement [20]. Accordingly, we have two types of rules: *bottom-up* rules and *top-down* rules.

Bottom-up rules are used in scene interpretation. For instance, some rules could be defined to identify objects from track measures; i.e., to obtain SCOB instances from TREN instances. An example rule may be: create a person instance when an unidentified track larger than a predefined size is detected inside a region of the image.


Top-down rules create suggested action instances from the current interpretation of the scene, the historical data, and the predictions. These actions are retrieved and interpreted by the GTL through the GTL/CL interface, resulting (if not discarded) in corrections of the tracking parameters, modifications of the data structures managed by the algorithm, etc. Recommendations can be generated at different abstraction levels; in practice, that means that feedback rules have terms of the TREN, SCOB, ACTV, and IMPC ontologies in their antecedent; and terms of the RECO ontology in their consequent. An example rule may recommend ignoring a track associated to a person that is inside an area marked as a mirror. More details of the use of RACER abductive rules in our framework can be found in [21].

4 Context-based Information Retrieval: Establishing CI from Knowledge Sources

Context-based Information Retrieval (CBIR) [22] is an independent, pre-processing step that occurs before a process that requires contextual information. A general CBIR procedure operates by examining an input, such as sensor data. It uses that input to constrain the knowledge that is available to the reasoner. This process is depicted in Fig. 6.



(a) Initial mark-up of the scenario



(b) Additional instances added during the processing

Figure 5. Correspondence between Scene Entities and Ontology Instances

² <http://www.racer-systems.com/>




Figure 6. General CBIR Processing Pipeline

The CBIR procedure receives an input that contains the contextual constraints and other information about the situation, and information from background knowledge sources (BKS) containing any domain knowledge that will be evaluated by the CBIR procedure. In the counterinsurgency domain these BKS can include:

- **Ontologies:** Ontological information about the domain entities (e.g., places, people), categorical relationships among them, and rules governing those relationships.
- **Previous Situations:** Information stored from previously process constraints.
- **Axioms:** Additional axioms for reasoning about the domain, including relationships not reasoned about in ontologies (e.g., part-hood)

With the contextual constraints provided and access to the BKS the CBIR procedure produces a subset of the BKS, called the contextual information (CI), to be added to the contextual constraints for use by a reasoning engine capable of performing forward-inference. The reasoning engine then performs forward-inferencing to either; (1) infer new information, called enhanced information, from the contextual information, which can be used to provide additional information about fusion estimates (a priori), or (2) infer a contradiction between the contextual constraints and the CI (a posteriori).

The CBIR routine is important not only because it accomplishes the above tasks in the soft fusion process, but also because it constrains the amount of information used from the BKS. Though CBIR routines can vary in implementation, they will all return a subset of the BKS and most will only examine a portion of the BKS using heuristic search methods. There are many large-scale data sources that are useful in the counterinsurgency domain, such as the NGA GEOnet Names Server³, and these heuristics can save computational time complexity when running as part of a typical fusion pipeline that needs additional information from such sources. Below one use of CBIR in counterinsurgency is presented.

4.1 Contextual Enhancement of Counterinsurgency Message Data

Work on the Soft Target Exploitation and Fusion (STEF) research project [23] has resulted in a message repository that contains soft-information from a manufactured counterinsurgency scenario. A message in this domain is a timestamped, small English paragraph containing information about activity that might be of interest to a counterinsurgency operative. An example message is as follows:

02/10/07 American contractor in Yarmuk said ambulance service is deteriorating; he told of a friend who suffered a heart attack in Qahtan Square, a short distance south of the hospital. His friend had to wait nearly an hour before the ambulance arrived.

Messages like the above contain information about various entities in the counterinsurgency domain (e.g., Yarmuk). Since these entities may have more information about them contained among the various BKS and such information might be useful to an Army field operative, a method for associating this additional information with the entities is desired. To solve this problem CBIR can be used in a process called *contextual enhancement*.

The contextual enhancement process operates by using a graph representation of the message, which uses labeled nodes to represent the entities in the message and labeled arcs between those nodes to represent the relationships specified in the message. The CBIR procedure uses the nodes in this graph as contextual constraints, and operates as described previously. The enhanced information is then provided to the Army field operative, along with the original graph representation of the message, in order to make tactical decisions.

5. Conclusions

Modern applications of information fusion require the employment and exploitation of all available information due to the complexities of these problems and the frequent difficulties of achieving sufficiently-rich observational data. Contextual information can often help in the formation of fused state estimates that are improvements over those developed from observational data alone, but of course such exploitation comes at a price of complexity that has to be factored into the cost-effectiveness of any prototype design. In this paper, we have reviewed the basic design frameworks for such contextual exploitation strategies, and given examples of some of our prototype applications that show that workable and effective designs can be realized. However, much remains to be done in defining and designing more robust and efficient approaches for use of contextual information; the fusion community needs to caucus on this issue and develop holistic research strategies that will result in unified design patterns for contextual exploitation.

³ <http://earth-info.nga.mil/gns/html/index.html>

Acknowledgements

The UC3M Team gratefully acknowledge that 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 DPS2008-07029-C02-02.

UC3M also thanks Prof. James Llinas for his helpful comments during his stay, which has been supported by the collaboration agreement ‘Chairs of Excellence’ between University Carlos III and Banco Santander.

The US/UB Team gratefully acknowledge that this research activity is supported by a Multidisciplinary University Research Initiative (MURI) grant (Number W911NF-09-1-0392) for “Unified Research on Network-based Hard/Soft Information Fusion”, issued by the US Army Research Office (ARO) under the program management of Dr. John Lavery

References

- [1] A. M. Fosbury, T. Singh, J. L. Crassidis, and C. Springen, *Ground target tracking using terrain information*, 10th Int. Conference on Information Fusion, 2007, Quebec City, Canada, pp. 1-8.
- [2] M. Kandefer, and S. C. Shapiro, *A categorization of contextual constraints*, Biologically Inspired Cognitive Architectures: Papers from the AAAI Fall Symposium, 2008, Technical Report FS-08-04, AAAI Press, Menlo Park, CA, USA, pp. 88-93.
- [3] A. Dey, and G. Abowd. *Towards a Better Understanding of Context and Context-Awareness*, CHI Workshop on the What, Who, Where, When, and How of Context-Awareness, The Hague, Netherlands, 2000.
- [4] H. Ekbja, A. Maguitman. *Context and Relevance: A Pragmatic Approach. Modeling and Using Context*, CONTEXT 2001, Dundee, UK, Lectures Notes in Artificial Intelligence 2116, pp. 156-169.
- [5] J. Llinas, C. Bowman, G. Rogova, A. Steinberg, E. Waltz, and F. White. *Revisiting the JDL data fusion model II*, 7th Int. Conference on Information Fusion, 2004, Stockholm, Sweden, pp. 1218-1230.
- [6] C. Nowak. *On ontologies for high-level information fusion*, 6th Int. Conference on Information Fusion, 2003, vol. 1, Cairns, Australia, pp. 657-664.
- [7] R. Arndt, R. Troncy, S. Staab, L. Hardman, and M. Vacura. *COMM: Designing a well-founded Multimedia Ontology for the Web*. 6th Int. Semantic Web Conference, 2008, Busan, South Korea, 30-43.
- [8] W. Lee, T. Bürger, and F. Sasaki. *Use Cases and Requirements for Ontology and API for Media Object 1.0*, Online, W3C Working Draft, 2009, URL <http://www.w3.org/TR/media-annot-reqs/>.
- [9] A. R. François, R. Nevatia, J. Hobbs, R. C. Bolles, and J. R. Smith. *VERL: an ontology framework for representing and annotating video events*, IEEE Multimedia, Vol. 12, No. 4, pp. 76-86, 2005.
- [10] U. Westermann, and R. Jain. *Toward a Common Event Model for Multimedia Applications*, IEE Multimedia, Vol. 14, No. 1, pp. 19-29, 2007.
- [11] M. Kokar, and J. Wang. *Using ontologies for recognition: An example*, 5th Int. Conference on Information Fusion, vol. 2, 2002, Annapolis, MD, USA, pp. 1324-1330.
- [12] L. Snidaro, and G. L. Foresti. *Knowledge representation for ambient security*, Expert Systems, Vol 24, No. 5, pp. 321-333, 2007.
- [13] B. Neumann, and R. Möller. *On Scene Interpretation with Description Logics*, Image and Vision Computing, Vol. 26, pp. 82-101, 2008.
- [14] M. M. Kokar, C. J. Matheus, and K. Baclawski. *Ontology-based situation awareness*, Information Fusion, Vol. 10, No. 1, pp. 83-98, 2009.
- [15] C.S. Regazzoni, V. Ramesh, and G.L. Foresti, *Scanning the Issue/Technology*, Special Issue on Video Communications, Processing, and Understanding for Third Generation Surveillance Systems, Proceedings of the IEEE, Vol. 89, No. 10, pp. 1355-1367, 2001.
- [16] M. Valera, and S.A. Velastin. *Intelligent distributed surveillance systems: a review*, IEE Proceedings - Vision, Image, and Signal Processing, Vol. 152, No. 2, pp. 192-204, 2005.
- [17] F. Bremond, and M. Thonnat, *A context representation for surveillance systems*, 4th European Conference on Computer Vision, 1996, Workshop on Conceptual Descriptions from Images, Cambridge, UK.
- [18] J. Gómez-Romero, M.A. Patricio, J. García, J.M. Molina. *Ontological representation of context knowledge for visual data fusion*, 12th Int. Conference on Information Fusion, 2009, Seattle, WA, USA, pp. 2136-2143.
- [19] C. Elsenbroich, O. Kutz, and U. Sattler. *A case for abductive reasoning over ontologies*, OWL workshop: Experiences and Directions, 2006, Athens, GA, USA.
- [20] J. Gómez-Romero, M.A. Patricio, J. García, and J.M. Molina. *Context-based reasoning using ontologies to adapt visual tracking in surveillance*, 6th IEEE Int. Conference on Advanced Video, 2009, Genoa, Italy, pp. 226-231.
- [21] J. Gómez-Romero, M.A. Patricio, J. García, and J.M. Molina. *Towards the implementation of an ontology-based reasoning system for visual information fusion*, 3rd Annual Skövde Workshop on Information Fusion Topics, 2009, Skövde, Sweden.
- [22] M. Kandefer, and S. C. Shapiro. *An F-Measure for context-based information retrieval*. 9th Int. Symposium on Logical Formalizations of Commonsense Reasoning, 2009, Toronto, CA, USA, pp. 79-84.
- [23] K. Sambhoos, J. Llinas, and E. Little. *Graphical methods for real-time fusion and estimation with soft message data*, 11th Int. Conference on Information Fusion, 2008, Cologne, Germany, pp. 1-8.