University of Huddersfield Repository

Olszewska, Joanna Isabelle and McCluskey, Thomas L.

Ontology-Coupled Active Contours for Dynamic Video Scene Understanding

Original Citation

Olszewska, Joanna Isabelle and McCluskey, Thomas L. (2011) Ontology-Coupled Active Contours for Dynamic Video Scene Understanding. In: Proceedings of the 15th IEEE International Conference on Intelligent Engineering Systems (INES),. IEEE, Poprad, High Tatras, Slovakia, pp. 369-374. ISBN 978-1-4244-8954-1

This version is available at http://eprints.hud.ac.uk/id/eprint/10829/

http://eprints.hud.ac.uk/

Ontology-Coupled Active Contours
for Dynamic Video Scene Understanding

Joanna I. Olszewska and Thomas L. McCluskey

School of Computing and Engineering, University of Huddersfield
Queensgate, Huddersfield, HD1 3DH, United Kingdom
{j.olszewska, t.l.mccluskey}@hud.ac.uk

Abstract—In this paper, we present an innovative approach coupling active contours with an ontological representation of knowledge, in order to understand scenes acquired by a moving camera and containing multiple non-rigid objects evolving over space and time. The developed active contours enable both segmentation and tracking of multiple targets in each captured scene over a video sequence with unknown camera calibration. Hence, this active contour technique provides information on the objects of interest as well as on parts of them (e.g. shape and position), and contains simultaneously low-level characteristics such as intensity or color features. The ontology we propose consists of concepts whose hierarchical levels map the granularity of the studied scene and of a set of inter- and intra-object spatial and temporal relations defined for this framework, object and sub-object characteristics e.g. shape, and visual concepts like color. The system obtained by coupling this ontology with active contours can study dynamic scenes at different levels of granularity, both numerically and semantically characterize each scene and its components i.e. objects of interest, and reason about spatiotemporal relations between them or parts of them. This resulting knowledge-based vision system was demonstrated on real-world video sequences containing multiple mobile highly-deformable objects.

I. INTRODUCTION

Automatic understanding of dynamic video scenes is a challenging task as it requires both an efficient and robust technique to capture visual information and its appropriate representation in terms of semantic concepts and relations.

In this work, we developed a knowledge-based vision system enabling the automatic analysis of the more general case of dynamic real-world scenes acquired by a moving camera with an unknown calibration and an arbitrary image resolution. The studied scenes typically contain multiple non-rigid objects potentially interacting with each other and quickly evolving on a clutter and changing background.

Vision systems dealing with the extraction and the tracking of video objects or parts of them usually represent the targets as elliptical centroids [6] or minimum bounding rectangles (MBR) [4] and most of the time they are restricted to specific conditions such as static backgrounds. In contrast, the method we used in this work is based on active contours [10] that are deformable two-dimensional curves evolving in the scene image plane from an initial position to the foreground boundaries. The adopted approach [16] can accurately extract non-rigid objects of interest without any restrictive assumption about the camera movement or target motion and can feed our system with visual information at different levels of scene granularity such as entire objects as well as parts of them and their related low-level features.

In order to interpret scenes, some recent works have developed ontologies for establishing a semantic structure of image visual content. The notion of ontology was defined in computer science by [7] as an explicit specification of a conceptualization. The term is borrowed from Philosophy and it refers to the subject of existence. In Artificial Intelligence (AI), an ontology is a formal specification of knowledge and is constituted by a specific vocabulary used to describe a certain reality and by a set of explicit assumptions regarding the intended meaning of this vocabulary. The ontological approach has been widely applied to various applications requiring knowledge representation, e.g. in air traffic control [13]. In Computer Vision, the proposed conceptualizations of visual scenes could refer to low-level features [5], objects [12], spatial relations between objects [9] or are event-based and thus define some specific domain ontologies [1]. Hence, each of these approaches is only partially covering the entire range of general semantic concepts necessary for the interpretation of dynamic scenes captured by mobile cameras and most of them are only dedicated to the study of scenes captured by stationary cameras.

Recently, some other works propose to integrate semantic modeling into vision systems e.g. by means of automatic labeling [11] or based on Bayesian networks [2]. However, these methods do not consider spatial relationship between objects such as occlusions.

In our presented system, active contours are coupled with a spatio-temporal visual ontology we designed. On one hand, active contour enables the system to gather both numeric properties of the objects of interest such as position, area, etc. and visual information (e.g. shape, color of parts of the objects of interest). On the other hand, the semantics of the scene are introduced in the system through the spatio-temporal visual ontology we have defined as well as spatiotemporal relations between and within the multiple objects detected with the active contour technique to support reasoning about scenes and parts of them.

The original contributions of this presented work are:
- the architecture of a new knowledge-based vision system enabling a general and effective approach for automatic understanding of dynamic scenes captured by mobile camera with multiple mobile highly-deformable objects of interest;
- the visual-related knowledge representation in an ontological form in which hierarchical levels of the concepts strictly map the different levels of granularity of the
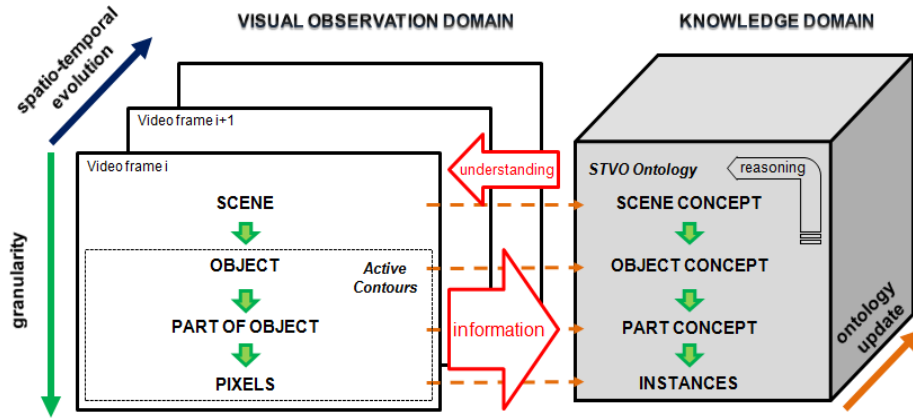
Fig. 1. Architecture of our system coupling Active Contours with the Spatio-Temporal Visual Ontology.

scene, leading to an efficient scalability and usability;

- the new definitions of inter- and intra-object spatial relative directional relations in regards to the *o'clock* notion, bringing semantic meaningful concepts and effective interpretation of visual information.

The paper is structured as follows. In Section II, we present our ontology-coupled active contour approach. The resulting knowledge-based vision system, that aims to automatically understand dynamic scenes, has been successfully tested on real-world standard video sequences as reported in Section III and discussed in Section IV. Conclusions are drawn up in Section V.

II. ONTOLOGY-COUPLED ACTIVE CONTOURS

Our approach depicted in Fig. 1 consists of coupling multi-target active contours with the developed spatio-temporal visual ontology. Concepts of this ontology map the granularity of the visual scene. The ontology contains also inter- and intra-object spatiotemporal relations defined for this framework and used for reasoning purpose.

We introduce in Section II-A how multi-target active contours could automatically provide visual information about the scene or parts of it. Then we describe in Section II-B our ontology designed for the conceptualization of visual dynamic scenes captured by a moving camera with unknown properties. Next we present in Section II-C the architecture of the resulting knowledge-based vision system conceived for the automatic understanding of dynamic scenes.

A. Multi-Target Active Contours (MTAC)

To automatically gather visual information (e.g. shape) about objects of interest and their parts as well as numeric data such as target position or area, we use the active contour method [10].

One advantage of this technique is its intrinsical robustness towards target's appearance changes that occur in dynamic scenes because of deformations of the non-rigid objects and/or movements of the camera.

On the other hand, when efficiently implemented, active contours enable both segmentation and tracking of highly-deformable object of interest. For this purpose, we have

adopted the approach proposed by [16] which provides in a computationally effective way shape and position of the target object and allows the extraction of related information such as intensity and color values e.g. in the (R,G,B) space.

Hence, based on this parametric active contour approach, the detection and extraction of a video object and its associated visual and geometrical information could be done as follows. Given an initial contour, the active contour algorithm is applied on a video frame. Once it reaches its final state, the active contour accurately delineates the object of interest. Moreover, since feature information such as intensity and/or color are extracted during the active contour process, they remain available for all the system.

In this paper, we extend this approach to the extraction of multiple objects of interest. In that case, the initialization is done separately for each target. The algorithm is then processed for each object, in a parallel way. The resulting active contours are thus quickly generated. We observe that the computed active contours (MTAC), as shown e.g. in Figs. 4 (a) and (e), are highly robust and do not present split-and-merge issues. For all these reasons and due to the separate management of each contour, the corresponding target objects could be precisely extracted and parts of the objects could be identified as well.

B. Spatio-Temporal Visual Ontology (STVO)

STVO ontology was designed by using a top-down approach [3], in the way its hierarchical concepts specifications map the granularity of the observed visual scene. Its structure could be seen on Fig. 2. Here the depth of the ontological tree is five, corresponding to visual levels of video sequence, scene, object, parts of objects and pixels, respectively. This architecture enables a higher computational efficiency in terms of navigability than deeper ontologies e.g. [12], whereas STVO covers a complete range of concepts necessary for the interpretation of dynamic scenes captured by mobile cameras.

To develop STVO domain knowledge, we have chosen Protégé v4 [8] and we have adopted the OWL standard language [15] to express the complex semantic entities and their relations. Its description logic (DL) allows us to perform
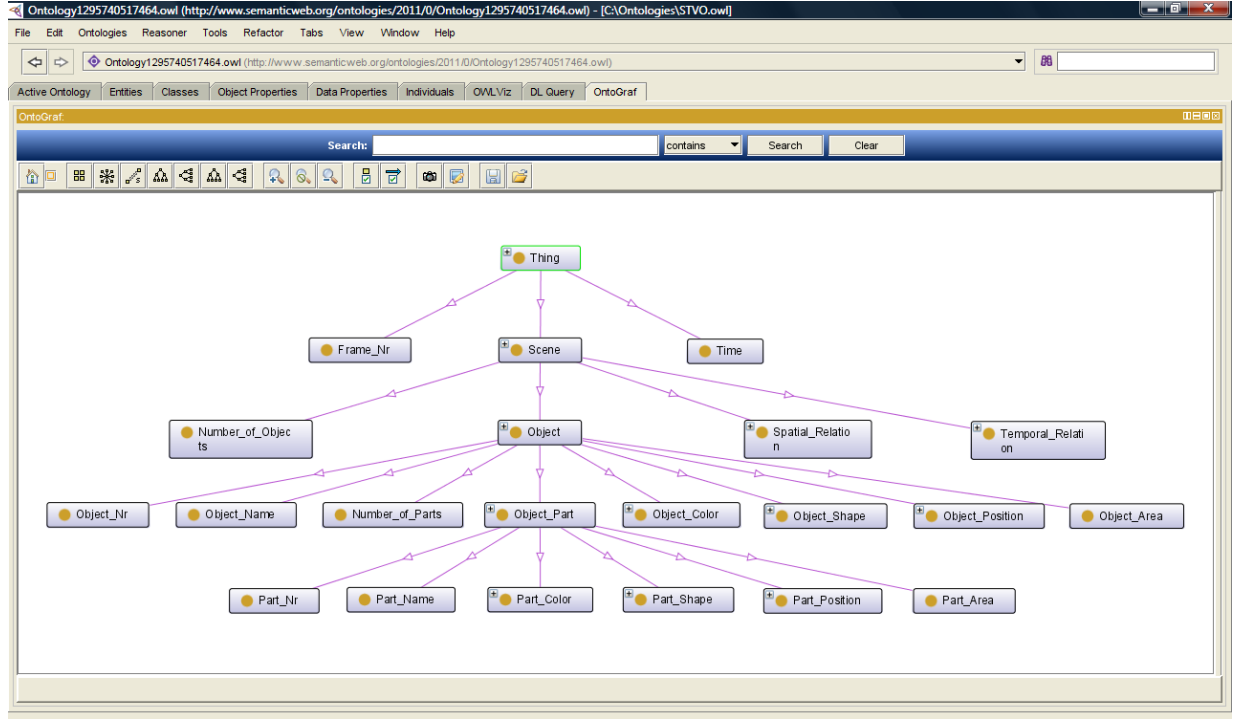
Fig. 2. Spatio-Temporal Visual Ontology (STVO).

automated reasoning on it and to formulate queries related to both semantic and numerical properties of visual scenes.

In the remaining part of this Section, we present the semantic concepts defined in our STVO ontology as well as the relations between them defined in the framework of dynamic scenes captured by a mobile camera.

A scene could be considered as an aggregate of objects of interest all linked by spatio-temporal, geometrical and photometric relations.

Hence, in our ontology, an object of interest is conceptualized in terms of shape, position and other attributes mapping the visual observation domain properties. Features such as texture are not relevant in STVO because of the blurring effect of the camera movement. Moreover, as we assume the general case of the unknown calibration of the mobile camera, properties such as object motion are not considered.

For the temporal relations, we focus on the main notions *hasAppeared*, *hasDisappeared*, and *isInScene*. These first two relations are related to the *start* and *end* ones [14], [20].

The DL description of the STVO property *isInScene* is as follows:

$$
\begin{aligned}
isInScene \sqsubseteq\ & Temporal\_Relation \\
& \sqcap hasAppeared \\
& \sqcap \neg hasDisappeared
\end{aligned}
\tag{1}
$$

where the object properties are defined through the existence or not of the *Object_Position* in a scene.

In STVO, we define spatial relations between and within objects that we have categorize in three families: (i) the topological relations, (ii) the directional relative positions, and (iii) the relative distances, sizes, and areas.

For (i), we adopt the RCC-8 model [19] consisting of eight basic relations between two objects: disconnected (DC), externally connected (EC), equal (EQ), partially overlapping (PO), tangential proper part (TPP), tangential proper part inverse (TPPi), non-tangential proper part (NTPP), non-tangential proper part inverse (NTPPi).

For (ii), we introduce in our ontology new semantically meaningful concepts in order to provide a unified framework for inter- and intra-object spatial relations and to reduce the uncertainty [9] on the relative positions between two objects. Our new formal specification of the relative directional relations is based on the clock space that numbers the scene plane as a clock's face (see Fig. 3 (a)). Indeed, in real-world crowded scenes, saying that an object is at the right of another could certainly be not enough discriminant as it could be many of them. Indicating the o'clock relative position could reduce the uncertainty related to the directional relative positions among objects as demonstrated in Section III.

As an example of how we have formalized the *o'clock* concept, we define the *2 o'clock* concept in DL as follows. Let O_{REF} be the object of reference, and let the object we want to know its relative directional relation with be called O_{REL}. Let $Angle$ be the relative angle between the line O_{REF}–O_{REL} and the X axis of the image plane. Then:

$$
\begin{aligned}
isAt2clockOf \sqsubseteq\ & Spatial\_Relation \\
& \sqcap Relative\_Distance\_Relation \\
& \sqcap \exists O_{REF} \\
& \sqcap \exists O_{REL} \\
& \sqcap \exists Angle2clock \\
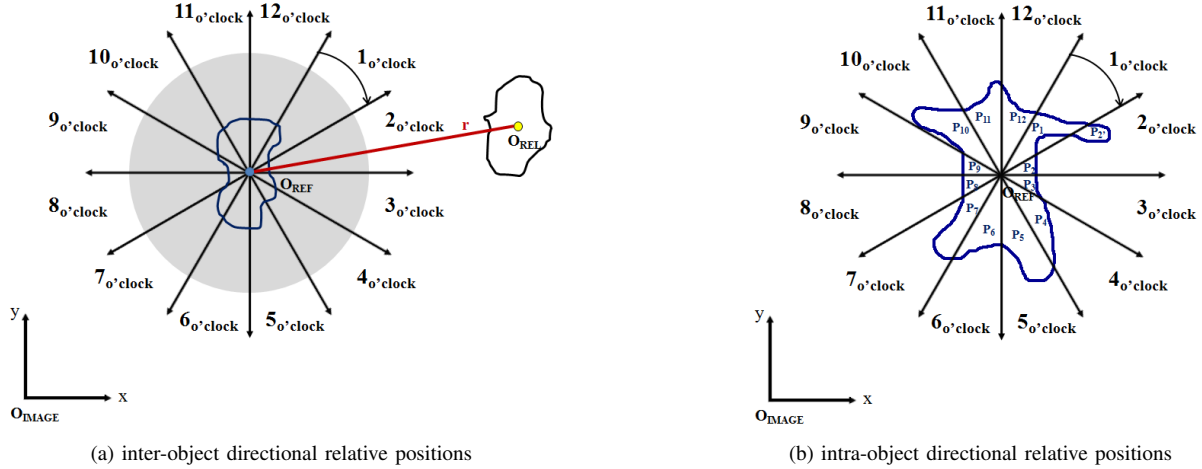& \sqcap \exists inverse.isAt8clockOf
\end{aligned}
\tag{2}
$$

(a) inter-object directional relative positions

(b) intra-object directional relative positions

Fig. 3.   Proposed o'clock-semantic meaning of directional relative positions.

with

$$Angle2clock \equiv Angle$$
$$\sqcap \exists angle.value_{\leq \frac{\pi}{6}} \tag{3}$$
$$\sqcap \exists angle.value_{>0}$$

Equation (3) denotes the set of angles which have values lower than or equal to $\pi/6$ and higher than 0. In this example, $_{\leq \frac{\pi}{6}}$ is a predicate over the real number domain $\mathbb{R}$.

The *2 o'clock* concept has the inverse property *8 o'clock*. Indeed, if a relative object (O_{REL}) is at *2 o'clock* of a reference object (O_{REF}), the reference object is at *8 o'clock* of the relative object (O_{REL}).

Moreover, the traditional relative intra-object relations (*Above_Of*, *Below_Of*, *Left_Of*, *Right_Of*) are embedded in our concept. Indeed, e.g. the property *isAt3clockOf* could be seen as equivalent to the traditional directional *Right_Of* relation.

Semantically meaningful spatial relative relations between p parts P_i of an object are introduced for the first time in an ontology in our work and we called them intra-object relations as shown on Fig. 3 (b). We formalize them using our new *o'clock* concept which does not induce an arbitrary division of the target but a partition taking automatically into account object's concavities and convexities, leading e.g. to 13 parts as in Fig. 3 (b).

For (iii), we define the concept of being *close* or *far* in terms of a proportion between r, the distance in the image plane between O_{REF} and O_{REL} and the size of the width and height of the scene image. We also propose the semantic concept of relative area and size of objects. These notions could help for example in scene understanding in situations such as an object is on the first plane, or there is a close-up on a target object.

In our STVO ontology, the color concept is defined by 16 basic color keywords defined in [21] as well as by its related explicit (R,G,B) value in the red (R), green (G), blue (B) color space. In the experiments carried out on real-world scenes acquired by mobile camera, this choice gave satisfactory results as shown in Section III and ensures computational effectiveness. However, these semantic concept values could be extended to more keywords if any application requires that.

In this work, we define the concept of the color of an object (*Object_Color*) containing p parts P_i $with$ $(i = 1, ..., p)$ as follows:

$$Object\_Color \equiv Part\_Color.P_1$$
$$\sqcup ... \sqcup Part\_Color.P_p \tag{4}$$

with

$$Part\_Color \equiv Color$$
$$\sqcap \exists hasRChannelValue_{=R_p}$$
$$\sqcap \exists hasGChannelValue_{=G_p} \tag{5}$$
$$\sqcap \exists hasBChannelValue_{=B_p}$$

In this example, $= R_p$, $= G_p$, and $= B_p$ are each a predicate over the number domain $[0, 255]$. The DL definition of the concept of the object color proposed in Eq. (4) takes into account the color of the different parts and not necessarily the average value over the whole detected object. Hence, this novel definition is more discriminant for objects with inhomogeneous color properties as it occurs in real-world scenes.

C. Resulting Knowledge-Based Vision System

In our resulting system (see Fig. 1), the active contour method automatically extracts from frames composing the visual observation domain the information to feed the knowledge domain. This is performed by mapping the numerical and visual properties with the semantic concept representations and relations. Hence the STVO ontology is updated. Since many of STVO relations take the form of logic rules on which automatic reasoning can take place, this leads to the automatic understanding of complex dynamic scenes.

Results of the analysis of scenes with multiple target objects interacting with each other are presented in Section III.

(a) frame 831 - active contour results   (b) ontology - query 1   (c) ontology - query 2   (d) ontology - query 3

(e) frame 1216 - active contour results   (f) ontology - query 4   (g) ontology - query 5   (h) ontology - query 6
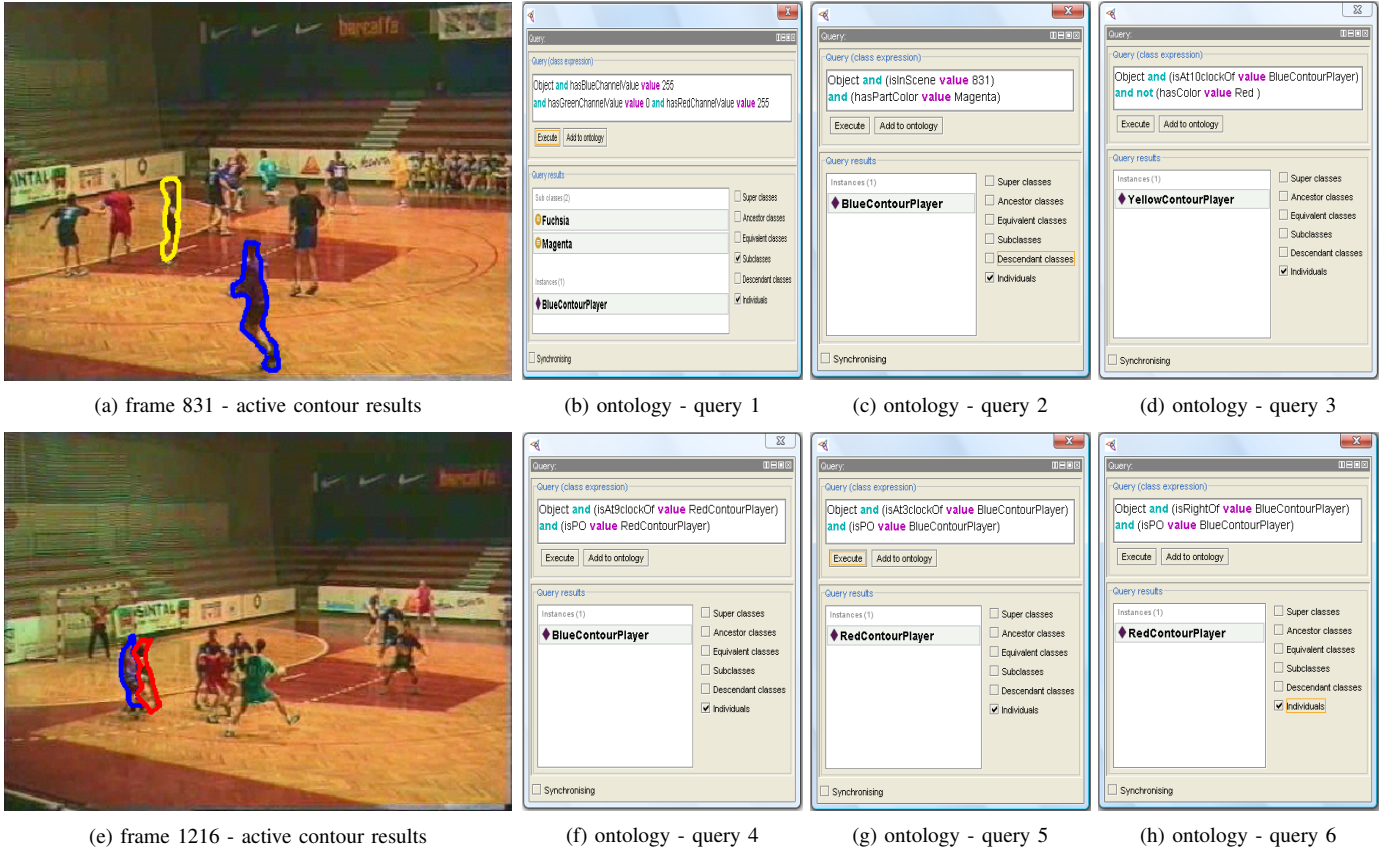
Fig. 4. Results of our method tested for several target-objects in a real-world handball video sequence. First column: visual observation domain. Second to forth column: knowledge domain. First row: frame 831. Second row: frame 1216.

III. EXPERIMENTS

Our ontology-coupled active contour approach has been validated on several real-world video sequences from well-known standard datasets: CVBASE, SCEPTRE, and BEHAVE. In this paper, experiments have been carried out on the CVBASE handball match [18], where the dynamic scenes are acquired by mobile cameras and where players are considered as the objects of interest. This video is chosen because of difficulties for tracking/segmentation (e.g. highly-cluttered background, moving camera, etc.) and for ontologies since crowd, mobile objects and video frames' low resolution challenge spatial, temporal, and appearance ontological relations, respectively.

In the first experiment, queries are focused on the color of the objects of interest and are related to the scene 831 displayed on Fig. 4 (a). In the first case (see Fig. 4 (b)), the system takes directly into account the numerical (R,G,B) values mapped from the visual scene to the ontology. After reasoning, the system is able to provide the related semantic meaning i.e. the color name and could find the object owning this color attribute in the scene. In the second case (see Fig. 4 (c)), the same query is specified in a semantic way. Both answers report that it is the player delineated by the blue contour who owns some magenta-color properties, that corresponds to the information associated with database.

In the second experiment (see Fig. 4 (d)), the third query shows the importance of the *o'clock* concept as well as of the use of multi-attribute queries. This gives a better discrimination of the objects, and leads to unique answers being obtained. In fact, six objects of interest are left of the blue-contour player. Asking about how many objects are at *10 o'clock* reduces this set of possible objects to two. Adding a question about the color attribute e.g. expressed in a negative way enables the system to find a unique solution i.e. to localize the object delineated by the yellow contour in the visual domain based on semantic questions.

In the third experiment, queries are related to the directional relative positions between objects of interest on scene 1216 shown on Fig. 4 (e). Queries four and five illustrate the inverse properties of the *o'clock* concept. Indeed, on Fig. 4 (f), the query is about the object at *9 o'clock* of the red-contour player that partially occluded him; the answer is that it is the blue-contour player. On Fig. 4 (g), it is the inverse query (about the object at *3 o'clock* of the blue-contour player) and thus the system provides the right answer i.e. the inverse one of the previous case (it is the red-contour player). The sixth query (see Fig. 4 (h)) illustrates the equivalence between the meaning of *isAt3clockOf* and *isRightOf* relations.

All these results are shown on Fig. 4 (a)-(h) and all the carried out experiments demonstrate that on one hand, our system is able to automatically understand the visual scenes of the match and on the other hand, our system outperforms the state-of-the-art ones as discussed in Section IV.

IV. Evaluation and Discussion

The performance of our ontology-coupled active contours have been evaluated for the whole proposed system as well as for each of its parts.

To quantitatively assess the detection system, we define the $precision\ of\ detection = \#\,correct\ contours\,/\,\#\,targets$. For Fig. 4 (e), the accuracy values of the targets' detection by the different techniques are presented in Table I.

frame 1216	[10]	[17]	MTAC
precision of detection	0 [%]	50 [%]	100 [%]

TABLE I
ACCURACY RESULTS FOR THE TARGET DETECTION

Regarding our targets' detection subsystem (MTAC), running multiple active contours in parallel is in fact a challenging solution not previously proposed in the literature. Indeed, it is only possible (see Table I) if based on the recent, fast and robust active contours proposed by [16]. Hence, for Fig. 4 (e), traditional active contours fail. The level-sets [17] initialized for both players are merged in a single active contour, while snakes [10] collapse. Furthermore, our approach for multi-target active contours outperforms state-of-art techniques of segmentation (usually suffering from over-segmentation, thus not providing clear foreground/background segmentation necessary for our application) and of tracking (often needing static cameras as in [9], [12] or characterizing objects with centroids [6] that include the foreground and parts of the background, and then not offering a sufficient object representation in terms of shape and/or appearance in our case).

Regarding our STVO ontology, its design outperforms the state-of-art ones because it enables faster search time because our tree depth is 5 and not 8 like in [12]. On the other hand, our system allows not only simple queries like in state-of-art ontologies based on only one type of relations such as [9], but also complex queries based on multiple criteria (spatial, temporal, shape, appearance).

To quantitatively assess the visual ontology, we define the $precision\ of\ answers = \#\,correct\ answers\,/\,\#\,queries$. For the queries presented in Figs. 4 (b)-(d) and Figs. (f)-(h), the total accuracy values of the answers based on the different ontologies are presented in Table II.

Fig. 4	[9]	[12]	STVO
precision of answers	50 [%]	50 [%]	100 [%]

TABLE II
ACCURACY RESULTS FOR THE QUERIES

As shown in Table II, combinations of our semantically-meaningful relations provide significant improvements by reducing the set of the possible answers enabling more precise answers than those obtained using state-of-art relations. State-of-art systems such as [9], [12] fail on the queries and frames of Fig. 4 mainly because their higher ambiguity due to the use of only one type of relations, their lack of temporal relations and their less efficient visual system.

Regarding our overall system, it outperforms traditional systems used for scene understanding because, in addition of an efficient segmentation and tracking of multiple objects in video captured by static or mobile cameras, it provides *o'clock* inter-object relations outperforming state-of-art spatial relations (right, left, etc.) like in [9] because of a reduction of ambiguity e.g. in crowded situations (Fig. 4 (d)). Our system also includes *o'clock* intra-object relations outperforming state-of-art spatial relations because none specifically proposed in the literature and leading to a better characterization of the color semantic definition (set of colors of all parts of the object) than the the state-of-art definitions (single average value for complex objects like a player with shorts, shirt, etc. of different colors [5]). Thus, our ontology-coupled active contour system allows a better objects' description thanks to the availability of information on both whole and parts of object rather than just on parts such as tracking-by-parts state-of-art techniques like [4].

V. Conclusions

In this paper we have presented (a) an extended parametric active contour method which copes with multiple targets, and provides information about the visual observation domain; (b) a spatio-temporal visual ontology that provides a granular conceptualization of the visual observation domain as well as a set of novel inter- and intra-object spatio-temporal relations defined in this work in the 2D space. The designed system architecture consists in coupling (a) with (b), leading to scene knowledge which can be automatically reasoned with within description logic.

References

[1] M. Bertini, A. Del Bimbo, G. Serra, C. Torniai, R. Cucchiara, C. Grana, and R. Vezzani, "Dynamic pictorially enriched ontologies for digital video libraries", *IEEE Multimedia*, vol. 16, no. 2, pp. 42-51, April 2009.

[2] H. Cheng and R. Wang, "Semantic modelling of natural scenes based on contextual Bayesian networks", *Pattern Recognition*, vol. 43, no. 12, pp. 4042-4054, December 2010.

[3] O. Corcho, M. Fernandez-Lopez, and A. Gomez-Perez, "Methodologies, tools and languages for building ontologies. Where is their meeting point?", *Data and Knowledge Engineering*, vol. 46, no. 1, pp. 41-64, July 2003.

[4] E. Corvée and F. Brémond, "Body parts detection for people tracking using trees of histogram of oriented gradient descriptors", in *Proceedings of the IEEE International Conference on Advanced Video and Signal-Based Surveillance*, August 2010, pp. 469-475.

[5] S. Dasiopoulou, V. Mezaris, I. Kompatsiaris, V. K. Papastathis, and M. G. Strintzis, "Knowledge-assisted semantic video object detection", *IEEE Transactions on Circuits and Systems for Video Technology, Special Issue on Analysis and Understanding for Video Adaptation*, vol. 15, no. 10, pp. 1210-1224, October 2005.

[6] J. Gomez-Romero, M. A. Patricio, J. Garcia, and J. M. Molina, "Context-based reasoning using ontologies to adapt visual tracking in surveillance", in *Proceedings of the IEEE International Conference on Advanced Video and Signal-Based Surveillance*, September 2009, pp. 226-231.

[7] Thomas R. Gruber, "Towards principles for the design of ontologies used for knowledge sharing", *International Journal of Human-Computer Studies*, vol. 43, no. 5-6, pp. 907-928, November 1995.

[8] Matthew Horridge, *A Practical Guide to Building OWL Ontologies Using Protégé 4 and CO-ODE Tools, Ed. 1.2*, University of Manchester, 2009.

[9] C. Hudelot, J. Atif, and I. Bloch, "Fuzzy spatial relation ontology for image interpretation", *Fuzzy Sets and Systems*, vol. 159, no. 15, pp. 1929-1951, August 2008.

[10] M. Kass, A. Witkin, and D. Terzopoulos, "Snakes: Active contour models", *International Journal of Computer Vision*, vol. 1, no. 4, pp. 321-331, January 1988.

[11] L.-J. Li, R. Socher, and L. Fei-Fei, "Towards total scene understanding: classification, annotation and segmentation in an automatic framework", in *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, June 2009, pp. 2036-2043.

[12] N. Maillot and M. Thonnat, "Ontology based complex object recognition", *Image and Vision Computing*, vol. 26, no. 1, pp. 102-113, July 2008.

[13] T. L. McCluskey and M. West, "Ontological Engineering and ATC", *Air Traffic Technology International*, pp. 63-65, 2009.

[14] G. Nagypal and B. Motik, "A fuzzy model for representing uncertain, subjective and vague temporal knowledge in ontologies", in *Proceedings of the International Conference on Ontologies, Databases and Applications of Semantics*, November 2003, pp. 906-923.

[15] B. Neumann and R. Moeller, "On scene interpretation with description logics", *Image and Vision Computing*, vol. 26, no. 1, pp. 114-126, January 2008.

[16] Joanna Olszewska, *Unified Framework for Multi-Feature Active Contours*, PhD Thesis, UCL, 2009.

[17] N. Paragios and R. Deriche, "Geodesic active contours and level sets for the detection and tracking of moving objects", *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 22, no. 3, pp. 266-280, March 2000.

[18] J. Pers, M. Bon, and G. Vuckovic, Computer Vision Based Analysis in Sport Environments (CVBASE) Dataset, Team handball subset. Available online: http://vision.fe.uni-lj.si/cvbase06/dataset.html, May 2006.

[19] D. A. Randell, Z. Cui, and A. G. Cohn, "A spatial logic based on regions and connection", in *Proceedings of the International Conference on Knowledge Representation and Reasoning*, October 1992, pp. 165-176.

[20] J. G. Stell and M. West, "A four-dimensionalist mereotopology", in *Proceedings of the International Conference on Formal Ontology in Information Systems*, November 2004, pp. 261-272.

[21] Scalable Vector Graphics (SVG) 1.1 (2nd Ed.), W3C TR, Recognized color keyword names. Available online: http://www.w3.org/TR/SVG/types.html#ColorKeywords, June 2010.