

# Semantically Enhanced UAVs to Increase the Aerial Scene Understanding

Danilo Cavalieri, Vincenzo Loia<sup>ID</sup>, Senior Member, IEEE, Alessia Saggese, Member, IEEE,  
Sabrina Senatore, Member, IEEE, and Mario Vento, Member, IEEE

1 **Abstract**—Visual tracking supported by unmanned aerial  
2 vehicles (UAVs) has generated a lot of interest in recent years,  
3 especially in application domains such as surveillance, search for  
4 missing persons and traffic monitoring. The major challenges  
5 in visual tracking with small UAVs arise in the form of target  
6 representation, target appearance change, target detection and  
7 localization in real time computation. Reliable target detection  
8 depends on factors such as occlusions, image noise, illumination  
9 and pose changes, or image blur that may compromise the object  
10 labeling. To mitigate these issues, this paper proposes a hybrid  
11 solution: along with the tracked objects, scenes are completely  
12 depicted by adding contextual information, i.e., data describing  
13 places, natural features, or in general points of interest. Each  
14 scenario indeed is semantically described by ontological state-  
15 ments that define the context and then, by inference, support the  
16 object tracking task in the object identification and labeling. The  
17 synergy between the tracking methods and semantic modeling  
18 can bridge the object labeling gap, enhancing the scene under-  
19 standing and awareness when alarming situations are discovered.  
20 Experimental results are promising and confirm the applicabil-  
21 ity of the proposed framework in supporting drones in object  
22 identification and critical situation detection tasks.

23 **Index Terms**—Mobile camera, semantic Web, situation  
24 awareness, unmanned aerial vehicles (UAVs), video tracking.

25

## I. INTRODUCTION

26 IN THE recent years, aerial surveillance is becoming cru-  
27 cial in many safety-critical application domains, such as  
28 fire detection, traffic congestion or accidents, etc. Unmanned  
29 aerial vehicles (UAVs) represent a clear, low cost reply to  
30 ground-plane surveillance systems, in order to recognize alert-  
31 ing situations. Although UAVs should guarantee rapid time  
32 of response, especially when considering a victim's mortal-  
33 ity and morbidity after a severe injury accident, at the same  
34 time, they should also potentially fly in uncomfortable weather

Manuscript received August 1, 2017; accepted September 17, 2017. This  
paper was recommended by Associate Editor Z. Li. (*Corresponding author:*  
*Vincenzo Loia*.)

D. Cavalieri, A. Saggese, S. Senatore, and M. Vento are with the  
Dipartimento di Ingegneria dell'Informazione ed Elettrica e Matematica  
Applicata, University of Salerno, 84084 Fisciano, Italy (e-mail: loia@unisa.it).

V. Loia is with the Dipartimento di Scienze Aziendali—Management and  
Innovation Systems, University of Salerno, 84084 Fisciano, Italy.

This paper has supplementary downloadable material available at  
<http://ieeexplore.ieee.org>, provided by the authors. This includes videos that  
show articulated situations set on the road with some critical events. This  
material is 34.1 MB in size.

Color versions of one or more of the figures in this paper are available  
online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TSMC.2017.2757462

conditions, that could be too dangerous for a manned aircraft. The main issue with the UAVs is the difficulty in acquiring a high-level description of the scenes appearing in the video sequence, only from the object detection, identification and tracking algorithms. Enabling a UAV to acquire a complete description of the scenario, during the flights and then, to recognize critical scenes from a video sequence is indeed, a very useful and desirable capability. To this purpose, a robust tracking to handle complex tasks such as object identification [1] and event detection [2] is required. Many studies focus indeed on alleviating common problems related to UAV video tracking such as camera resolution, camera shaking, illumination change and appearance change of the background [3], [4], but also to achieve optimal trajectory tracking [5] and alleviate vehicle routing problems [6]. Although many tracking algorithms can deal with occlusion, split objects, shadows, and reflection, object tracking suffers in object labeling [7]; moreover, camera movements add further problems to the object tracking algorithm. Most of the algorithms presented in literature work on object tracking with a fixed camera, and, on moving camera, the traditional background subtraction algorithms are not applicable. For this reason, most approaches concentrate on a single object class recognition task, for instance, pedestrians [8] or crowds [9]; vehicles [10], or their relations [11] and the main applications in this area converge on a single type of scenario [12]. In the light of these observations, our basic idea is to improve the object tracking task in a video sequence, augmenting the tracking data with the contextual data, i.e., complementary data related to the surrounding background objects of the scene, in order to acquire a more complete scene information to alleviate tracking issues. Data from tracking algorithms and additional background information are collected and coded into ontological statements: the role of the semantic Web technologies in the modeling of tracked objects and their relations with other objects in the environment is critical for object classification and labeling, especially when a moving camera is involved (videos came from a flying UAV).

The rationale behind the semantic enrichment of the scene description is indeed to exploit the contextual information from surrounding background objects such as places, buildings, rivers, and roads, in general, points of interest (POIs) to better identify and then label the tracked objects. Just to give an example: a tracked object on a road (where other similar tracks appear), more likely is a car; if a similar track appears on a river, it will be a boat. An object trajectory (hereinafter

81 a *track*) whose predominant color is red could be a fire: if  
82 it appears in a wood, then probably it represents a dangerous  
83 situation; but if it appears on a beach, it could be just a bonfire,  
84 which is likely not dangerous.

85 Semantically coded scenes feed a knowledge base, which  
86 becomes a source to query and to infer comprehensive, high-  
87 level information about the objects enclosed in the scene and in  
88 the video. The ontology-based modeling of scene from video  
89 sequences can enhance the video tracking methods, supporting  
90 the object labeling and providing a high-level interpretation of  
91 the scenes in the video sequence: objects are discovered and  
92 automatically labeled with the actual name; at the same time,  
93 event and object interactions in the scene can be monitored  
94 so that a critical situation can be detected when an alarming  
95 event (pedestrian on the road, car crash, fire, etc.) is revealed.  
96 This paper presents a novel approach which exploits the syn-  
97 ergy between the tracking methods and semantic technologies  
98 to bridge the object labeling gap, enhance situation aware-  
99 ness, as well as detect and classify alerting situations. The  
100 UAV, provided with a camera mounted on board, can recog-  
101 nize moving and background objects, that populate the scene,  
102 and relations/interactions between them. The semantic tech-  
103 nologies provide the way to collect all these data and produce  
104 a comprehensive description of the scene, that will be used  
105 to infer additional information. The UAV becomes “aware” of  
106 the situations occurring in the evolving scenario and, from  
107 the contextual (background) data, can also individuate and  
108 interpret alerting situations. In a nutshell, the proposed study  
109 presents the following.

- 110 1) An approach that merges data from video tracking  
111 output (by a mobile camera) with high-level environ-  
112 mental data, in order to provide a straightforward scene  
113 description and understanding.
- 114 2) A semantically enhanced knowledge base which collects  
115 data describing dynamic scenarios involving various  
116 kinds of objects and environments.
- 117 3) A reasoning process to produce additional contextual  
118 knowledge (feeding the knowledge base), useful to  
119 provide object classification and event detection.

120 This paper is organized as follows. Section II provides a  
121 brief overview about the main works on video tracking, and  
122 the role of semantic technologies in supporting object track-  
123 ing and labeling approaches. Section III sketches the overall  
124 framework model; then Section IV introduces the semantic  
125 model on which the framework lays; specifically, a closer  
126 look at the semantic modules and their interactions is given  
127 in Section V. Experiments in Section VI show the effective-  
128 ness of this approach in the object classification and labeling  
129 as well as the scene understanding. The conclusion and future  
130 work close this paper.

## 131 II. RELATED WORKS

132 Although high-level knowledge-based scene understanding  
133 from moving camera is a hot topic for improving decision-  
134 making processes and supporting video-tracking activities like  
135 moving object detection and tracking [1], [2], [14], there  
136 are a few related works in literature studying the problem.

137 This depends on the fact that the moving camera causes  
138 a lack of reference points in the scene which affects both  
139 object detection and tracking activities, and high-level scene  
140 modeling. Then, most of studies focus only on low-level data  
141 coming from video or at most adding only few environmental  
142 variables to the problem. Moreover, *a priori* knowledge based  
143 on static context is not suitable with a dynamic environment  
144 like a scene taken from a moving drone with an on-board cam-  
145 era, which can record many different environments with many  
146 different moving object kinds moving in it.

147 The moving camera adds new problems to object detection  
148 and tracking, especially because there is no fixed background,  
149 which makes the distinction between self moving objects and  
150 environmental elements more difficult [15]–[18]. Therefore,  
151 consolidated fixed camera techniques like background sub-  
152 traction [19] cannot work because environmental element  
153 pixel data change with the moving camera. To constrain this  
154 problem, studies on the moving camera video-tracking make  
155 assumptions on the environment and camera; for instance,  
156 they assume *a priori* that the environment is finite and well  
157 known [12], [20], [21]; camera movements are constant or con-  
158 strained; tracking is carried out on only one object [8], [10].  
159 Some studies also achieve object recognition by object clas-  
160 sification in predefined classes, even though many issues as  
161 low resolution [22], motion blur [23], and prohibitive camera  
162 shots [24], [25] need to be addressed.

163 The goal of our approach is not a new object tracking  
164 algorithm but a knowledge-based model to understand scene  
165 dynamics at a high level, and to better improve object classi-  
166 fication, as well as to support complex decision tasks. For  
167 scene understanding, many works propose pattern recogni-  
168 tion methods to recognize scene elements or regions [26].  
169 Classification results are quite limited and do not provide  
170 a deeper and high-level understanding of the scene, which  
171 is required when dealing with evolving scenarios. Generally,  
172 most of these methods work exclusively on low-level pixel-  
173 based data, such as color, shape, and position. For a better  
174 contextual awareness of the scene, it is necessary to retrieve  
175 and model the environmental data. Environmental data can be  
176 used to build an initial knowledge on environment, useful to  
177 contextualize features and the moving objects. Studies [27]  
178 have proposed approaches that build spatio-temporal contexts  
179 exclusively on the tracking data.

180 In order to get a deeper knowledge of the environment,  
181 this paper presents a framework adopting semantic techniques  
182 to model a dynamic environment starting from some basic  
183 features. Semantic technologies allow building a high-level  
184 description of the environment and its elements, based on het-  
185 erogeneous information gathered from different sources. They  
186 provide a machine-oriented representation of the scenario and  
187 the situations evolving in the scenario. It is the main role, along  
188 with the inference process that, applied to the built model,  
189 can enhance the knowledge about the evolving scenario. The  
190 enhanced knowledge is useful to understand complex situa-  
191 tions in the current scenario, even though it is far from  
192 the comprehensive prediction of possible dangerous situations,  
193 especially when unpredictable behaviors happen, nor it is able  
194 to quantify the chance of an event occurring [28]. A synergistic

195 approach that exploits consolidated methodologies (for example, deep learning-based methods) could alleviate the issues  
 196 related to the foreseeing of future unexpected/unpredictable  
 197 events.  
 198

199 In the light of these observations, our idea is to exploit  
 200 the semantic technologies to provide a semantic enhancement  
 201 in the object labeling and scene understanding, in order to  
 202 suggest critical situations (i.e., situations where for example, the spatial relations among objects are out of allowed  
 203 range) and eventually, to wisely support a decision. Semantic  
 204 technologies are also used in video-tracking but with fixed  
 205 camera. They are mainly used for data fusion of low-level  
 206 data coming from different sensors [29], and for data fusion  
 207 between low-level data and contextual variables [30]. The  
 208 main goal of methods proposed in literature is to alleviate  
 209 tracking problems like occlusion, grouping, shadowing,  
 210 etc. [30]–[33]. They are also used to support object detection  
 211 proposing semantic segmentation techniques to classify pixel  
 212 regions in predefined classes [34]. Semantic segmentation  
 213 often involves deep learning techniques to classify pixels in  
 214 predetermined categories. These methods report good perfor-  
 215 mances on detecting environmental features [35], even though  
 216 they require many preacquired training samples [36]. In [30],  
 217 a framework producing high-level knowledge on an environ-  
 218 ment is presented. The application recognizes a door, a person  
 219 by dimensions and the action of the person entering in the  
 220 scene by the door. This approach needs to acquire the scene  
 221 *a priori* (e.g., door presence) with a fixed camera filming a  
 222 static and well-known environment. Our approach, instead, is  
 223 aimed at building an adaptable framework for various possi-  
 224 ble scenarios: it models at the semantic level, first, basic and  
 225 general concepts, suitable for every kind of scenario, and then,  
 226 employs a map-based tool to retrieve more specific environ-  
 227 mental data, to enrich the knowledge about the scenario. To  
 228 the best of our knowledge, there are no studies on the moving  
 229 camera video tracking employing semantic Web technologies  
 230 to improve the accuracy in the object identification. Moreover,  
 231 this hybrid approach can overcome the missing data problem  
 232 in tracking algorithms, and, thanks to the expressive power of  
 233 the semantics, produces a high-level description of events and  
 234 objects in the scene.  
 235

### 236 III. FRAMEWORK OVERVIEW

237 Fig. 1 shows the high-level scheme of the framework, with  
 238 all the components and their main interactions. The core of the  
 239 framework is represented by the semantic modules that, in the  
 240 figure, are enclosed in a black border square. This framework  
 241 extends a preliminary work, presented in [37].

242 The main input is a video recorded by a flying drone with  
 243 an on-board installed camera (top left of Fig. 1). The recorded  
 244 video is taken as input by the *tracking module* which extracts  
 245 the trajectories of the objects moving in the scene, frame by  
 246 frame. For each frame, tracked object dimensions and speeds  
 247 are also calculated. The other input of the framework is envi-  
 248 ronmental data (top right in Fig. 1): it is composed of specific  
 249 places called POIs, which are fixed geo-referenced points or areas  
 250 retrieved with Google Maps service, lying in the area where

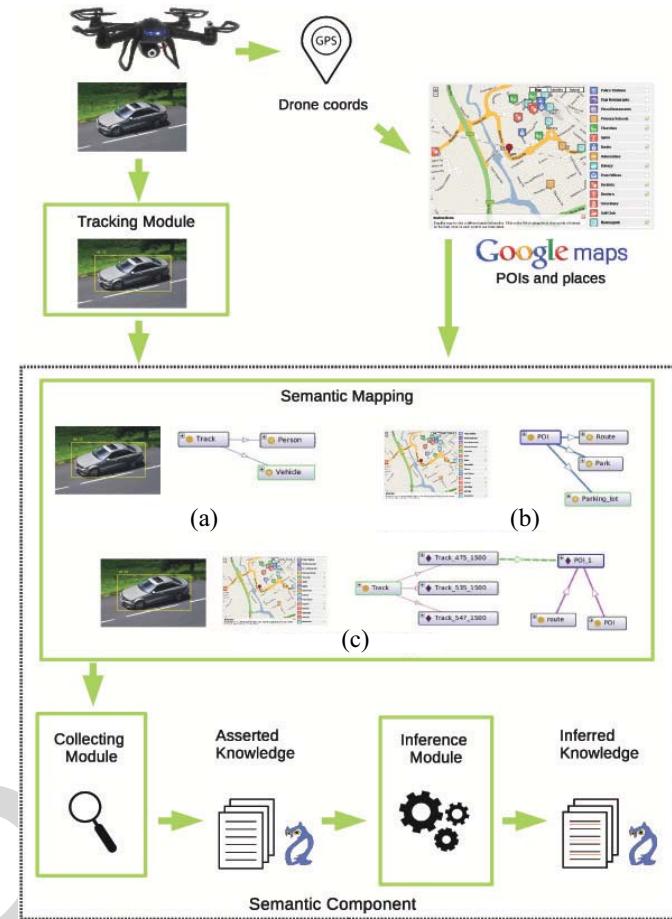


Fig. 1. Logical overview of the framework.

the drone flies. The video sequence and the object trajectories, 251 as well as the POIs retrieved with Google Maps are passed to 252 the *semantic mapping* module. This module in turn, translates 253 tracking and contextual data in semantic statements accord- 254 ing to TrackPOI ontology, an ad-hoc designed ontology 255 to model the on-the-road scenarios. The *semantic mapping* 256 module is composed of three submodules. The first one is *track* 257 *semantic mapper* which maps moving objects and frame data 258 in assertions about their identity, real dimensions, speed and 259 position. The *POI semantic mapper* aims at defining ass- 260 sumptions on the POIs data retrieved with Google Maps query. The 261 third module is *relation semantic mapper*: from the knowl- 262 edge base produced by the track and POI semantic mappers, 263 it extracts positional relations between tracked objects, and 264 tracked objects and POIs in the scene, and then generates the 265 corresponding assertions which feed the knowledge base. 266

The collected knowledge on tracks, POIs, and their rela- 267 tions is passed to a *collecting module*, which collects and 268 refines the contexts in the evolving scene. 269 Finally, this knowledge is passed to the *inference module*, 270 which deducts new assertions on tracks, POIs and relations, to 271 feed a comprehensive knowledge base, for a deep high-level 272 scene understanding. 273

Before detailing the framework component, a brief descrip- 275 tion of the ontology and its role in the semantic description 276 of main components is given in the next section. 277

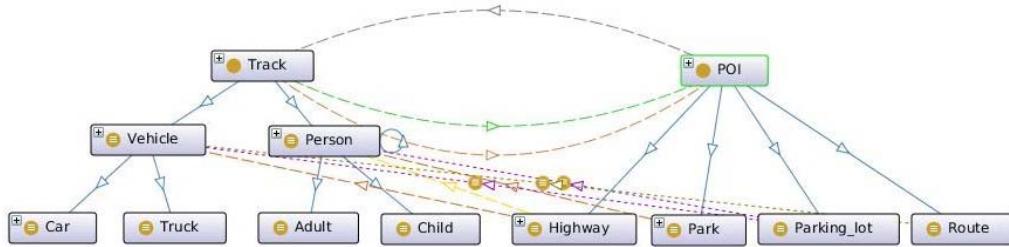


Fig. 2. TrackPOI ontology schema.

#### IV. ONTOLOGY DESIGN

The ontology design provides an explicit, shareable way to describe a domain of interest. In this paper, it describes geo-contextualized scenarios coming from video sequences, recognizing not just the single entities appearing but interpreting the whole evolving scene.

Fig. 2 shows TrackPOI ontology schema. The ontology presents two main classes *Track* and *POI*, which, respectively, model track and POI entities. *Track* class presents two subclasses: 1) *Person* and 2) *Vehicle*. These subclasses are the possible types of moving objects present in the scene and each one of them has other subclasses, for *Vehicle*: car, motorcycle, and truck; and for *Person*: adult and child.

POIs are structured data about urban and natural environments (e.g., roads, bridges, squares, buildings, stores, institutions, mountains, oceans, rivers, etc.) and localized by their own GPS position. In Fig. 2, only the subclasses *Highway*, *Park*, *Parking\_lot*, and *Route* are shown. Besides the class-subclass relation, transversal relations are also shown; a high-level generic relation *hasRelationWith* between *Track* and *POI* is given: it states that some relationship exists between some track and some POI. Specialized relations can be defined increasing the subclass depth, where well-designed relations link subclasses; for instance a specialized relation *isInTheAreaOf* exists between classes *Person* and *Park*. Further details about the designed relations are given in the following sections.

#### V. DEEP OUTLOOK ON SEMANTIC MODULES

Although the section presents all the components involved in the framework modeling, the main focus is given to the modules embedding semantic processes.

##### A. Tracking Module

The tracking module extracts the trajectories of those objects moving in the scene, and provides to the *semantic mapping* module the position (and the dimensions), frame by frame, of each object moving in the scene. The adopted tracking algorithm was proposed in [38]. It is based on three steps.

1) *Camera compensation*, which aims at estimating the movement of the camera so as to extract a kind of background of the scene representing the scene with no objects inside.

- 2) *Detection*, which extracts the position of the objects moving in the scene at the current instant of time.
- 3) *Tracking*, which aims at associating the objects detected at the current frame with the ones tracked up to the previous frames, so as to draw, for each object, their trajectories frame by frame.

In cases of a low probability of this association, a feed-forward approach is exploited: the parameters adopted during the detection step are indeed adaptively and automatically adjusted, so as to increase the probability of the association. The tracking algorithm is able to effectively deal with both merging (two or more objects identified as a single one during the detection step) and splitting (a single object partitioned into two or more parts during the detection step). Let us remark that the tracking algorithm has been just used to get tracked objects in a format suitable to be processed by our semantic components; its performance analysis is out of the topic of this paper and is presented in [38].

##### B. POI Semantic Mapper

The POI semantic mapper processes POIs and places appearing in video scenes. POIs are fixed objects in the scene by definition, and for this reason they can be considered as reference points, which are very useful to understand movements of objects present in a moving camera recorded video. Furthermore, POIs play another important role in modeling knowledge about the video scenes. In this paper indeed, POI data are also used to define a context by restricting domain about the scenario. For example, if the system retrieves POI data about a public park which generally not allows vehicle transit, the main moving objects walking and standing in the park area will be people and pets.

Google API provides Google Maps Geocoding API<sup>1</sup> and Google Maps Places API<sup>2</sup> to localize, geo-refer and retrieve specific data about POIs. Therefore, POI semantic mapper adopts Geocoding API to make reverse geocoding by a simple query, which takes a pair of coordinates and returns a Json/XML file with a human-readable address and related data about the area which the pair of coordinates corresponds to. The main retrieved data are POI identity or type, administrative area level, postal code, street address, and GPS area and position. For each retrieved POI, the POI semantic mapper generates a new POI individual, i.e., an instance of class *POI*

<sup>1</sup><https://developers.google.com/maps/documentation/geocoding/intro>

<sup>2</sup><https://developers.google.com/places/>

(described by *TrackPOI* ontology, Fig. 2), and codes its own retrieved data in RDF triples. Precisely, the POI position and area are retrieved from *geoRSS* ontology,<sup>3</sup> a well-known geographical ontology that provides geospatial properties of POIs, and then, they are integrated into our ontology.

Since Geocoding API returns data about macro places like university, park, zoo, etc., Google Maps Places API has been queried to get information about small and more simple POIs, which are additional reference points which a track can also interact with. Google Maps Places API indeed returns a list of 97 different places (e.g., bank, bar, and park) and their related information in a similar way to Geocoding API.

Let us notice that a place is not associated with an area (like macro places) but just a position identified with GPS coordinates. Queries submitted to both Google APIs are based on the drone coordinates (associated with every frame of the recorded video), whereas the covering radius to retrieve places is based on the distance covered by the drone.

In a nutshell, the POI semantic mapper retrieves data about the POI location, identity and its related data which contribute to depict the context of a scenario. Then, each POI appearing in a frame is coded as a POI individual, i.e., a class instance of the *TrackPOI* ontology.

### 385 C. Track Semantic Mapper

The *Track* semantic mapper is in charge of converting the tracking output, provided by the tracking module, in ontological statements. A tracked object or simply a track is a dynamic object present in the video, that moves in the scene. A track can rapidly change from a moving to fixed state and viceversa in a few of frames. In most cases, tracks move autonomously in the environment, and they often are living beings like humans or animals, but they can also be not living things carried, pushed or driven by living beings like vehicles, balls, etc. The study focuses on various outdoor environments mainly populated by humans, animals and vehicles. The tracking output includes data about the tracks and frames which they are in. The *Track* semantic mapper reads the tracking output and extracts data about each track. For each track identified, the mapper generates a *Track* individual according to the proposed *TrackPOI* ontology (Fig. 2), where *Track* is the ontology class modeling the tracked objects. The *Track* individual describes a specific track in the frame, by means of the following properties: bounding box ID, track name, track real dimensions (width, height) in meters, track speed, track position (GPS coordinates), and frame ID which the track is related to. Another important added property is the *TrackPOI:hasRelationWith*, which relates the track individual to all the other tracks or POIs present in the same frame. All the generated *Track* individuals along with their own properties are added to the knowledge base.

### 412 D. Relation Semantic Mapper

The relation semantic mapper acquires the asserted knowledge on both tracks and POIs, generated by the other two

semantic modules, POI semantic mapper and *Track* semantic mapper, respectively. Its goal is to recognize relations between the distinct entities appearing in the frames by the analysis of the possible interactions between them in the evolving scene. The investigated relations are between a track and a POI, and between two tracks representing different entities. Relations are coded as asserted knowledge and are crucial to delineate the context of the scenario. The study focuses particularly on the positional relations, suitable to analyze how a track is positioned with respect to fixed POIs in the scene, and how a track is interacting with other tracks in the scene. The analyzed relations cover the main grammar prepositions of place: *front of*, *behind*, *near*, *between*, *in*, and *on*. Each preposition describes a geometric relation between GPS coordinates of tracks and POIs.

As a first task, relation semantic mapper creates a relation instance *TrackPOI:hasRelationWith* between each track and POI, present in the same frame, then it adds them to the asserted knowledge. The mapper processes track-POI relations frame by frame, in order to specialize these relations with respect to the contextual information. For each frame indeed, it analyzes every *TrackPOI:hasRelationWith* statement among track-track and track-POI filtering out irrelevant relations (i.e., on tracks without a specified position).

Geometric calculations are applied on each relation per frame, taking into account: GPS coordinates of the tracks and POIs, positional speed of a track in the current and previous frame, track and POI real dimensions and track directions. The geometric data allows relation semantic mapper to discern positional relations between entities in the scene, for example, a track in the area of a POI, a track/POI in proximity of another track/POI, etc. More specifically, relation semantic mapper can produce six different specializations of the predicate *TrackPOI:hasRelationWith*, that are modeled in the *TrackPOI* ontology. They are detailed as follows.

- 1) *TrackPOI:isInTheAreaOf*: This predicate is used to relate two entities, *x* and *y*. The statement *x TrackPOI:isInTheAreaOf y* [see Fig. 3(a)] is produced if the GPS coordinates of *x* (a track or a POI) lie within the area of *y* (a POI). Generally, the area of *y* is retrieved by the Google Maps query on POIs, that provides two pairs of coordinates which, respectively, represent the north-east and south-west zone bounds. When the retrieved POIs lack of these area bounds, the relation semantic mapper defines a covering radius, according to place type and video features, to delineate an area for the POI.
- 2) *TrackPOI:isNear*: This predicate relates two entities *x* and *y* which are close to each other, and, similarly to predicate *TrackPOI:isInTheAreaOf*, the mapper enables us to specify a covering radius value to define the concept of closeness [see Fig. 3(b)].
- 3) *TrackPOI:hasDirection*: This is a predicate that can be specialized, according to spatial relations between two entities. The mapper calculates the track trajectory (considering positions in successive frames) which is translated in cardinal points; an assertion *x TrackPOI:hasDirection p* is produced for every frame

<sup>3</sup>[http://www.georss.org/rdf\\_rss1.html](http://www.georss.org/rdf_rss1.html)

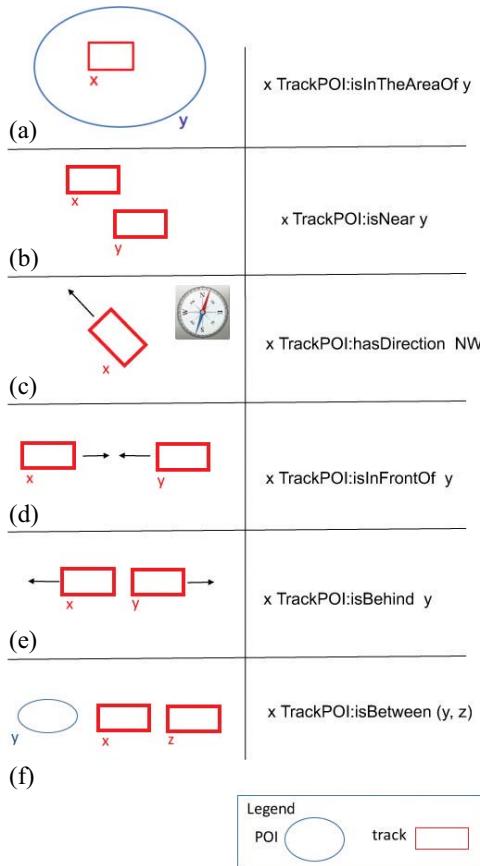


Fig. 3. Specializations of the predicate *TrackPOI:hasRelationWith*, according to the geometric relations among entities in the scene.

the track is in, where  $x$  is the track and  $p$  the cardinal point representing its direction [see Fig. 3(c)].

- 4) *TrackPOI:isInFrontOf*: Thanks to track direction per frame, the relation semantic mapper evaluates if two entities  $x$ , and  $y$ , are coming in front of one to the other. In detail, the assertion  $x$  *TrackPOI:isInFrontOf*  $y$  holds if the direction of  $x$  in a frame is opposite to the direction or position of  $y$ , where  $y$  is a POI and the distance between  $x$  and  $y$  in successive frames decreases [see Fig. 3(d)].
- 5) *TrackPOI:isBehind*: In a similar but reverse way, the mapper uses the predicate *TrackPOI:isBehind* to state  $x$  and  $y$  leave each other behind and proceed to opposite direction ways [see Fig. 3(e)].
- 6) *TrackPOI:isInBetween*: If the track  $x$  is between two entities  $y$  and  $z$ , relation semantic mapper can state that  $x$  *TrackPOI:isInBetween* ( $y, z$ ). This assertion holds if the position of  $x$  falls on the conjunction of  $y$  and  $z$  direction vectors, and if also the assertions  $x$  *TrackPOI:isNear*  $y$ ,  $x$  *TrackPOI:isNear*  $z$  hold [see Fig. 3(f)].

#### 493 E. Collecting Module

494 The collecting module completes the assertion process,  
 495 taking as input the asserted statements produced by the seman-  
 496 tic mappers. It synthesizes the semantic knowledge removing  
 497 redundant information, then checks if the statements are  
 498 consistent, according to *TrackPOI* ontology and drone data;

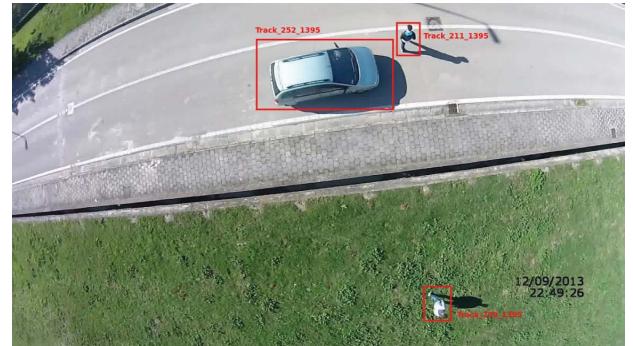


Fig. 4. Frame 1395 from video 1.

499 finally, statements are merged producing the definitive assen-  
 500 tional knowledge base (ABOX). The asserted statements  
 501 are conform to class and property definitions and restric-  
 502 tions of *TrackPOI* ontology (Fig. 2). The collecting module  
 503 checks if these restrictions are satisfied to guarantee a con-  
 504 sistent schema, especially on the relations between track  
 505 and track or track and POI. For example if a track  $x$   
 506 has a *TrackPOI:inFrontOf* relation with another track  $y$ ,  
 507 ( $x$  *TrackPOI:inFrontOf*  $y$ ),  $y$  must have the same relation with  
 508  $x$  ( $y$  *TrackPOI:inFrontOf*  $x$ ) because the *TrackPOI:inFrontOf*  
 509 property holds on two objects going toward each other.  
 510 *TrackPOI:inFrontOf* is indeed a symmetric property.

511 The statements produced are merged to form the definitive  
 512 asserted knowledge base on all the scenes of the video.

513 In order to give an example about the way collecting mod-  
 514 ule generates the comprehensive assertional knowledge, some  
 515 statements from a video sample are discussed. Fig. 4 shows  
 516 a video frame (whose identifier is 1395), showing two per-  
 517 sons: one walking on the grass, and the other one crossing the  
 518 road, in the proximity of a moving car. The frame shows the  
 519 corresponding two bounding boxes, identified by the track-  
 520 ing module. Algorithm 3 describes all the tracks appearing  
 521 in the frame by ontological statements (black rows 1–32 and  
 522 38–50). For the sake of simplicity, the statements in the form  
 523 of triples  $\langle \text{subject-predicate-object} \rangle$  are expressed in Turtle<sup>4</sup>  
 524 semantic language. Rows 6–18 show statements on the POI  
 525 entity named *POI\_1*: it is an individual of the *TrackPOI:POI*  
 526 class (row 8); it is a route (row 9). It is also described by asser-  
 527 tions on the area and the geo-position (rows 16 and 17), and  
 528 other Google Maps data like street address, country, admin-  
 529 istrative area level, and postal code (rows 10–15). Tracks  
 530 are described in black rows 20–32 and 38–50, precisely  
 531 *track\_211\_1395* and *track\_252\_1395* whereas the statements  
 532 about *track\_239\_1395* are not shown because very similar to  
 533 axioms about *track\_211\_1395*. The track objects are individ-  
 534 uals of the class type *Track* (rows 21–22 and 39–40) with  
 535 object dimensions and movements data like width, height,  
 536 and speed (rows 28–30 and 46–48) and GPS track position  
 537 (rows 31 and 49). The remaining assertions on track entities  
 538 are about movements and relations. Specifically, statements in  
 539 rows 23–27 and 41–45 describe the geometrical or positional  
 540 relations that the track has with other tracks and POIs in the

<sup>4</sup><https://www.w3.org/TR/turtle/>

---

```

1 Track
2 and ((hasRelationWith some (POI
3 and ((ObjectsAllowed some Person) or (ObjectsAllowed only Person)))
)
4 and (hasHeight only xsd:decimal[>= 100 , <= 220])
5 and (hasSpeed only xsd:decimal[>= 0 , <= 37])
6 and (hasWidth only xsd:decimal[>= 10 , <= 90]))
```

---

Listing 1. Equivalent class restriction for Person class.

same frame. Direction statements express the track direction from the previous frame with compass points (rows 32 and 50). The knowledge base about each video frame is passed to the inference module, which, by the analysis of the acquired statements, can infer new statements about the scene.

#### 546 F. Inference Module

547 The inference module implements the reasoning component  
 548 of the framework. As stated, it produces new axioms about the  
 549 scene objects and the context built on their relations, with the  
 550 aim of enriching the knowledge base and enhancing the scene  
 551 understanding. The new knowledge produced by this module  
 552 is aimed mainly at providing object classification and labeling,  
 553 as well as recognizing critical events in scene sequences. The  
 554 inference engine is built on OWL class equivalent, subclass  
 555 and disjoint restrictions as well as on rules implemented with  
 556 Semantic Web Rule Language (SWRL).<sup>5</sup>

557 Class restrictions are useful to provide precise modeling  
 558 of classes to guarantee a straightforward reasoning process in  
 559 inferring new assertions and provide their accurate classifi-  
 560 cation. Track individual classification has been designed on  
 561 *Track* class restrictions which involve the main bounding box  
 562 features and contextual data. Track subclasses (e.g., person,  
 563 vehicle, etc.) are defined as equivalent class restrictions, that  
 564 express a high-level definition of the class type (see Fig. 2).  
 565 Algorithm 1 shows an example of *Person* class modeling as a  
 566 class restriction: an individual of the class *Person* requires real  
 567 dimensions (width, height) falling in a specific range which  
 568 reflects the precise dimensions of a person seen from a top  
 569 view (rows 4 and 6), and a speed which is acceptable for a  
 570 moving human being (row 5). Relations with the context have  
 571 to be also specified: the *Person* class definition requires at  
 572 least a relation with a POI which admits persons in its area  
 573 (e.g., the presence of the person in a park). The admissibil-  
 574 ity about the presence of a *Track* individual in a POI area is  
 575 expressed by the *ObjectsAllowed* property (rows 2, 3). In order  
 576 to define the right relations between entities, allowable *Track*  
 577 individuals for a certain POI have to be specified. Algorithm 2  
 578 for instance, outlines that the only allowable *Track* types for  
 579 a *Park* class are individuals of the *Person* class: only person  
 580 tracks can appear in a park area. If an object cannot be related  
 581 to some POI (e.g., a person is not in a Park), the object can  
 582 be recognized by its dimensions and speed, but it is marked  
 583 as not recognized by context with a special property.

584 Restriction-based reasoning supports the object labeling, but  
 585 the expressive power is not suitable for situation understand-  
 586 ing. Rules have been designed to recognize situations and

<sup>5</sup><https://www.w3.org/Submission/SWRL/>

---

1 POI and (ObjectsAllowed only Person)

---

Listing 2. Equivalent class restriction for park class.

consequently alerting events occurring when restrictions on the scenes involving track objects and POIs do not hold. Each rule has been designed to verify that, in a specific situation, no unexpected event is revealed. When it is triggered, it identifies the critical event occurred on the involved objects and forces the system to provide an alert. For example, let us consider the following rule:

*Person(?x)  $\wedge$  Route(?y)  $\wedge$  Vehicle(?z)  $\wedge$  isInTheAreaOf(?x, ?y)*  
 *$\wedge$  isInTheAreaOf(?z, ?y)  $\wedge$  isNear(?x, ?z)  $\rightarrow$  isInDangerOn(?x, ?y).*

The SWRL rule describes an alerting situation that happens when a person is on a road. Given a *Route* individual *y*, and two tracks, respectively, representing a *Person* individual *x* and a *Vehicle* individual *z*, if *x* falls in the area of *y*, (*isInTheAreaOf*(?x, ?y)), the person *x* is in danger on the route *y*, especially because the vehicle *z* is coming (*isInTheAreaOf*(?z, ?y)  $\wedge$  *isNear*(?x, ?z)). The inferred property *isInDangerOn* represents an imminent alerting situation, a kind of prevision of a future dangerous situation. The reasoning process, mainly based on class restrictions and rules, works on one frame at a time. The inference module cycles on frames: it retrieves the statements related to it with a SPARQL query, and produces a subset of statements associated with each frame. Then, the inference module processes this statement subset and infers new statements, which are added to the knowledge base. Thus, when the inference module processes the statements for the frame in Fig. 4, new assertions are generated on the tracks *Track\_211\_1395* and *Track\_252\_1395*:

Algorithm 3 shows the final augmented knowledge, with an emphasis on the new assertions, red colored in Algorithm 3.

These assertions state that *Track\_211\_1395* is actually a person (row 35): so far the previous asserted statements (rows 21, 22) state that it was just a track and thanks to inference module, *Track\_211\_1395* is labeled as a person. *Track\_252\_1395* is a vehicle (row 52) and more specifically a car (row 53). Statements on the situation are also deducted: *Track\_211\_1395* is in an alerting situation, since it is on the route *POI\_1* (row 36).

## VI. EXPERIMENTAL RESULTS

An experimentation has been conducted on some videos, in order to evaluate our system performance in terms of object classification and capabilities in producing a complete scene understanding. The framework has been tested on full recorded videos, selecting consecutive sequences of frames from the start to the end of the video. Therefore, the framework acquires the tracking output about these frame sequences and processes it, building the knowledge base and running the reasoning process. In this way, the proposed framework is able to work on subsets of frame sequences at a time and returns the inferred new axioms within seconds. This feature reduces reasoning computation time and makes the whole system suitable for real-time applications (e.g., real-time video streaming).

```

1 @prefix trackpoi: <http://www.semanticweb.org/danilo/ontologies
2 /2016/1/ .
3 @prefix owl: <http://www.w3.org/2002/07/owl#> .
4 @prefix rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#> .
5 @prefix georss: <http://www.georss.org/georss/> .
6
7 # POI_1 triples
8 TrackPOI:POI_1 rdf:type owl:NamedIndividual ,
9 TrackPOI:POI_1 rdf:type TrackPOI:POI .
10 TrackPOI:POI_1 TrackPOI:Administrative_area_level_1 "Campania" .
11 TrackPOI:POI_1 TrackPOI:Administrative_area_level_2 "Provincia di
12 Salerno" .
13 TrackPOI:POI_1 TrackPOI:Administrative_area_level_3 "Fisciano" .
14 TrackPOI:POI_1 TrackPOI:Country "Italy" .
15 TrackPOI:POI_1 TrackPOI:Postal_code "84084" .
16 TrackPOI:POI_1 TrackPOI:Street_address "Anello Esterno, 84084
17 Fisciano SA, Italy" .
18 TrackPOI:POI_1 georss:box "40.775762,14.787031,40.772937,14.785706"
19
20 # track_211_1395 triples
21 TrackPOI:Track_211_1395 rdf:type owl:NamedIndividual .
22 TrackPOI:Track_211_1395 rdf:type TrackPOI:Track .
23 TrackPOI:Track_211_1395 TrackPOI:hasRelationWith TrackPOI:POI_1 .
24 TrackPOI:Track_211_1395 TrackPOI:hasRelationWith TrackPOI:
25 Track_239_1395 .
26 TrackPOI:Track_211_1395 TrackPOI:isNear TrackPOI:Track_252_1395 .
27 TrackPOI:Track_211_1395 TrackPOI:isInTheAreaOf TrackPOI:POI_1 .
28 TrackPOI:Track_211_1395 TrackPOI:hasHeight 58.309975411032354 .
29 TrackPOI:Track_211_1395 TrackPOI:hasSpeed 3.2 .
30 TrackPOI:Track_211_1395 TrackPOI:hasWidth 38.51992315031834 .
31 TrackPOI:Track_211_1395 georss:point "40.77454810242632,
32 14.78530388911672" .
33 TrackPOI:Track_211_1395 TrackPOI:hasDirection TrackPOI:N .
34 # inferred triples for track_211_1395
35 TrackPOI:Track_211_1395 rdf:type <https://schema.org/Person> .
36 TrackPOI:Track_211_1395 TrackPOI:isInDangerOn TrackPOI:POI_1 .
37
38 # track_252_1395 triples
39 TrackPOI:Track_252_1395 rdf:type owl:NamedIndividual .
40 TrackPOI:Track_252_1395 rdf:type TrackPOI:Track .
41 TrackPOI:Track_252_1395 TrackPOI:hasRelationWith TrackPOI:POI_1 .
42 TrackPOI:Track_252_1395 TrackPOI:hasRelationWith TrackPOI:
43 TrackPOI:Track_252_1395 TrackPOI:hasRelationWith TrackPOI:
44 TrackPOI:Track_252_1395 TrackPOI:isInTheAreaOf TrackPOI:POI_1 .
45 TrackPOI:Track_252_1395 TrackPOI:isNear TrackPOI:Track_211_1395 .
46 TrackPOI:Track_252_1395 TrackPOI:hasHeight 50.35861512770976 .
47 TrackPOI:Track_252_1395 TrackPOI:hasSpeed 3.4 .
48 TrackPOI:Track_252_1395 TrackPOI:hasWidth 13.428964034055936 .
49 TrackPOI:Track_252_1395 georss:point "40.77412995839289,
50 14.786379978047478" .
51 TrackPOI:Track_252_1395 TrackPOI:hasDirection TrackPOI:ENE .
52 # inferred triples for track_252_1395
53 TrackPOI:Track_252_1395 rdf:type <https://schema.org/Vehicle> .
54 TrackPOI:Track_252_1395 rdf:type TrackPOI:Car .

```

Listing 3. Assertional and inferred knowledge in Turtle statements for POI\_1, :track\_211\_1395 and :track\_252\_1395. The inferred triples are in red.

people or vehicles moving randomly in a scene. These videos<sup>6</sup> are recorded in the campus of the University of Salerno, specifically in an area which comprises a main route called Anello Esterno, heliport, parking, and some other places in the neighborhood like laboratories, department buildings, bars, etc. The drone mission has been planned with DJI Ground Station software, which provides the drone height and GPS data used by the framework in order to calculate object position, real dimensions, speed, as well as it retrieves the POIs appearing in the area. The first video shows two students walking together near the route until one goes away from the road and the other crosses the road while a car is coming. The second video captures the two students meeting once again, that are near and on a heliport while the helicopter is temporarily absent. videos 3 and 4 show a scenario similar to the one shown in video 1. These videos have been recorded between 3 and 4 P.M. in the summer. videos length is about a minute and half with a frame-rate of 25 frames/s. Each video has been divided into scenes of 15/22 s, time estimated for the scene evolution, where various events and objects appear. Additional videos, taken into account, are part of UAV123 dataset,<sup>7</sup> and present scenarios of people and vehicles moving in different environments such as route, roundabout, etc., even though their scenes do not present relevant alerting situations. videos are short (about 41 s), divided into three scenes, and present the same frame rate of 25 frames/s. Finally, the last two videos are retrieved from the Web and present more articulated situations, with some critical events. They seem interesting for the scenes evolution and the number of tracks appearing in the frame sequences. The experimentation is mainly based on the evaluation of the two main features: the capability of our system in the object identification and labeling and the recognition of critical or alerting situations during the scene evolution.

### B. Object Identification

Correct identification of objects depends strongly on tracking results, because the recognition is partly based on bounding box dimensions. Tests do not involve the evaluation of the tracking algorithm (studied in [38]), but the capabilities of the reasoning process to identify the object type. To this purpose, a handmade ground truth of the video, has been extended with a semantic annotation, describing the object type. Each bounding box has been annotated with a label that is the TrackPOI ontology class which the object (in the bounding box) belongs to. Class individuals returned by our framework are compared to the “semantic” ground truth annotations at frame level. The classes involved in the object classification are all subclasses of *Track* class, mainly: *Person* and *Vehicle*; also *Vehicle*’s subclasses *Car* and *Truck* (see Fig. 2) have been considered. The system evaluation in object identification is based on the precision and the recall measures. For each frame, all the tracks are retrieved with a SPARQL query, including the inferred statements about the class they belong to. According to the Pascal overlap criterion [39], if a calculated bounding box intersects with a bounding box in the ground truth for at

<sup>6</sup><https://drive.google.com/open?id=0B75yuWMeqbP5NVloZEIzc05jeW8>

<sup>7</sup><https://ivil.kaust.edu.sa/Pages/Dataset-UAV123.aspx>

TABLE I  
 OBJECT LABELING TEST RESULTS

Video	Scenes	Frames	Precision	Recall	F1 score
1	1	578	0.94	0.72	0.82
	2	578	0.83	0.75	0.79
	3	578	0.88	0.83	0.85
	4	579	1.0	0.44	0.61
2	1	421	0.83	0.81	0.82
	2	421	0.95	1.0	0.97
	3	422	0.66	1.0	0.80
	4	422	0.89	1.0	0.94
3	1	506	1.0	1.0	1.0
	2	505	1.0	1.0	1.0
	3	506	1.0	0.78	0.88
	4	505	1.0	0.86	0.92
4	1	375	1.0	1.0	1.0
	2	375	1.0	1.0	1.0
	3	375	0.96	0.84	0.90
	4	374	1.0	0.89	0.94
5	1	340	1.0	0.85	0.92
	2	340	0.82	1.0	0.90
	3	340	1.0	0.82	0.90
6	1	343	1.0	0.73	0.84
	2	344	1.0	1.0	1.0
	3	344	1.0	1.0	1.0
7	1	342	1.0	1.0	1.0
	2	342	0.94	1.0	0.97
	3	343	1.0	0.91	0.95
8	1	345	1.0	0.91	0.95
	2	345	0.97	1.0	0.98
	3	345	1.0	1.0	1.0
9	1	610	1.0	0.76	0.86
	2	610	1.0	0.87	0.93
	3	610	1.0	0.89	0.94
10	1	610	1.0	0.7	0.82
	2	610	1.0	0.87	0.93
	3	610	1.0	0.83	0.91

most the 50%, then the two class labels are compared. With this evaluation method, the comparison holds when one of the following cases occur.

- 1) The two labels are equal: our system has correctly recognized the moving objects, i.e., a true positive is returned.
- 2) The two labels are not equal: our system individuates and identifies the object but does not recognize it correctly; this is a case of a false positive (FP).
- 3) The bounding box returned by our system does not intersect with a ground truth bounding box: that means the system does not recognize a true and effective object; it represents a false negative (FN).

Table I shows results on object labeling, displaying precision, recall, and F1-score for each scene in a video. Ten videos have been analyzed, split in several scenes; each one composed of several frames (see *Frames* in Table I), varying in a range from 340 to 610. The frame rate is 25 for all the videos except the last two, with a frame rate equal to 30. Let us remark that the framework does not evaluate results at each frame because there are very little changes both in objects and events, in a second. For this reason, our framework analyses from 15 to 23 frames per scene.

The system performance generally can be considered satisfactory. Table I shows that the recall value is greater than

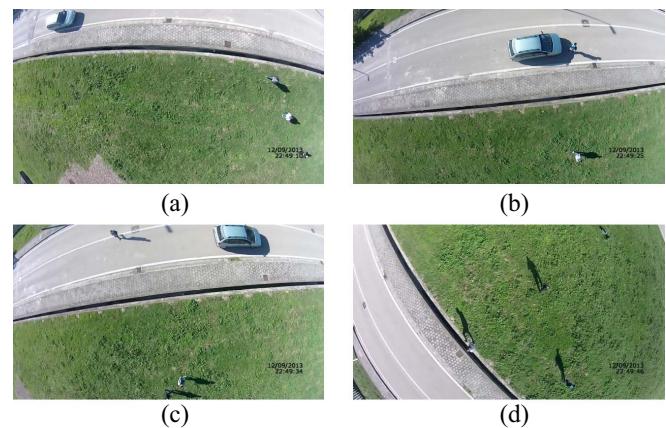


Fig. 5. Video 1, frames from scenes (a) 1, (b) 2, (c) 3, and (d) 4.

72% on all scenes of every video except for the first video, fourth scene, where the recall is 44%. Similar trend is observed for the precision, for each video scene. Fig. 7 shows the F1 score trend (and precision), by varying the recall. F1 curve has a strong growth with increasing recall starting from its lowest peak (0.61), due to low recall (0.44) on a scene, till to reach 1.0. On recall values greater than 0.7, there are just a few little variations due to some precision variations in value range 0.82–0.90. A precision decreasing also affects the F1 score curve, but the values are comprehensively high. The F1 trend reflects the good performance of the proposed system, whose capability at labeling objects improves with little improvements in tracking performance (growing recall values).

Looking at video 1, Fig. 5 shows four its own frames, taken respectively, from each scene. Video 1 recorded the two students walking close to the route, particularly, one of them crossing the road while a car is coming. In general, the system correctly recognizes the objects in scenes 1 and 3 [Fig. 5(a) and (c)]; and the precision for these scenes gets the highest values. Scene 2 [Fig. 5(b)] has some more FPs because the tracking algorithm returns some variable bounding boxes on the car and the student, when they are very close. The tracking algorithm does not detect the stopped car and generates some FNs. Scene 4 [Fig. 5(d)] presents the highest number of FN, because the tracking algorithm fails in the detection task, due to the ghost object presence as well as the environment element movements.

video 2 shows a different scenario: two students and other people walking and standing around to the area of the heliport. Four representative frames, each one taken from a different scene, are shown in Fig. 6. Also for this video, the system performances are good: the recall is always greater than 81% on all scenes, whereas the precision is generally high, except for the third scene where it reaches only 66%. In the first scene, three persons appear; the person walking to the center of the platform is always detected and correctly recognized, while the other two persons are not: when they are detected, their bounding boxes often have variable dimensions, with the resulting generation of possible FPs and FNs. Scene 2 shows the two students coming toward the center of the platform: they

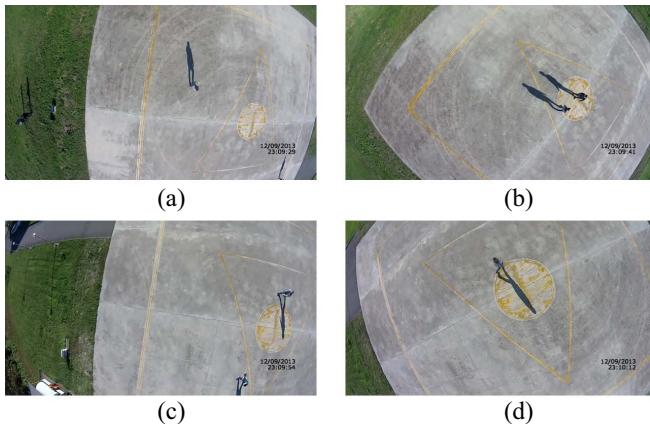


Fig. 6. Video 2, frames from scenes (a) 1, (b) 2, (c) 3, and (d) 4.

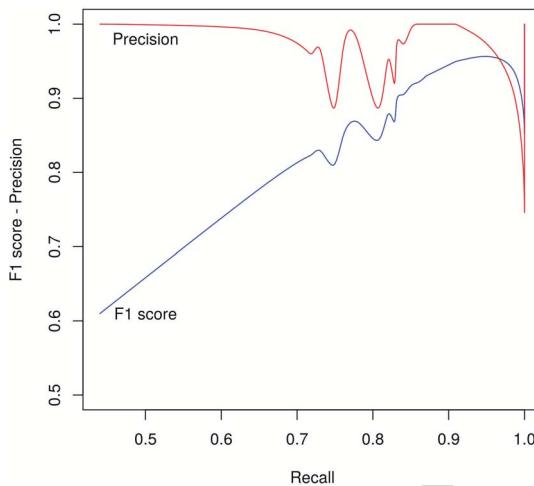


Fig. 7. Object labeling: F1-recall and precision-recall curves on the video collection.

768 are almost always correctly recognized except a very few FPs  
 769 in some frames. Scene 3 presents the two students walking  
 770 and stopping at the center of the heliport, in this case, the  
 771 object labeling gets more FPs, due to excessive variable bound-  
 772 ing boxes returned by the tracking algorithm. Finally, scene 4  
 773 shows students going back to the platform borders, and in this  
 774 case, the precision rises to 89% due to the presence of a few  
 775 FPs.

776 Videos 3 and 4 are so similar to video 1 in dynamics and  
 777 contents, but the tracking algorithm returns a better output  
 778 which is reflected in the results. In detail, in scenes 1 and 2 of  
 779 video 3, the system correctly recognizes all the people mov-  
 780 ing toward the route, while in scenes 3 and 4 there are some  
 781 FPs caused by the tracking algorithm output. For video 4, the  
 782 system perfectly recognizes all objects in the first two scenes,  
 783 while in scene 3, the tracking algorithm has some problems in  
 784 detecting the student crossing the route when he passes near  
 785 a car: this event badly affects the precision introducing with  
 786 some FPs.

787 Also for videos 5–8, taken from UAV123 dataset, Table I  
 788 presents good results on both precision and recall measures.  
 789 Video 5, first scene shows a man walking on the grass, then  
 790 he meets a second man (second scene) and finally in the

third scene, the couple stops and has a chat. The system cor-  
 791 rectly recognizes the two men in every scene, yet it gets poor  
 792 precision on the second scene, due to the presence of some  
 793 FPs caused by the excessive speed of the object movements.  
 794 Video 6 is set on a roundabout, the first and third scenes show  
 795 just a few cars while the second shows more cars going in  
 796 different directions. The framework recognizes perfectly all  
 797 the vehicles in every scene, as it is evident in the table; only  
 798 the recall for the first scene suffers the presence of some FNs.  
 799 videos 7 and 8 show similar route and roundabout scenarios to  
 800 those present in video 6, but taken from different illumination  
 801 conditions and camera angles. Precision and recall results are  
 802 quite similar to those of videos 5 and 6. Further tests have  
 803 been carried out on UAV123 dataset, by considering a sample  
 804 of 20 videos. They have been not reported in Table I because  
 805 present similar scene and tracked objects, and consequently,  
 806 similar results to videos 5–8.

807 Video 9 instead, captures an occurred car crash in the middle  
 808 of a crossroad. This leads to additional critical situations: the  
 809 crashed car forces the other vehicles to turn into the oncoming  
 810 traffic, furthermore, road crossing pedestrians are obscured by  
 811 the stopped vehicles. The video presents a great amount of  
 812 moving objects appearing in the scene. The system resulting  
 813 performance on this video are promising: precision values,  
 814 in particular, are around 100% on all scenes. Recall mean  
 815 value is around 80% with a minimum of 76%. Video 10 shows  
 816 a scenario similar to video 9. In this video there are more  
 817 cars and trucks, especially two long ones, that dangerously  
 818 swerve into the oncoming traffic. The results evidence a very  
 819 good precision: a great amount of objects are correctly labeled.  
 820 Recall has a minimum of 70% on the first scene where some  
 821 FNs occur, but on the other scenes gets very high values.

### C. Situation Understanding

823 The system performances are also measured in terms of  
 824 the situation detection, i.e., the individuation of critical events  
 825 occurring in the scene and involving detected objects. Our  
 826 system can detect events, by elaborating semantic data: it  
 827 processes specific statements produced by the reasoning pro-  
 828 cess (see Section V-F). These statements are rules or prop-  
 829 erties that can be activated when a critical/alerting situation  
 830 (described by some statements) is recognized (if the system  
 831 does not produce these statements, the scene is recognized as  
 832 safe). Rules are designed indeed on events involving Google  
 833 Maps POIs and moving objects, thus, when they are activated,  
 834 new statements are generated, describing situations classified  
 835 as critical.

836 In order to assess the capabilities of our system in detect-  
 837 ing critical situations, a ground truth of the critical events has  
 838 been manually built, and extended marking with annotation  
 839 each critical event occurred in every video scene. In videos 1–4  
 840 (described in the previous section), events that produce critical  
 841 situations are related to the road and the heliport. Specifically,  
 842 the events involve tracks, mainly persons and vehicles and  
 843 interact with both POIs. They are: 1) Ev1—walking or stand-  
 844 ing around the route or heliport and 2) Ev2—walking or  
 845 standing near (more close to) the route or heliport and, when

847 the POI is the route, the events are: 1) Ev3—crossing/running  
 848 the road when it is empty and 2) Ev4—crossing/running the  
 849 road when a vehicle is coming (valid for both persons and  
 850 vehicles), whereas for the heliport, they are: 1) Ev5—walking  
 851 or standing in the area of the heliport near the borders and  
 852 2) Ev6—walking or standing in the area of the heliport at  
 853 center.

854 videos 5–8 are set in route, park, and roundabout scenarios  
 855 showing a single kind of object (vehicle or man), thus they  
 856 present very few critical road-related events.

857 Events Ev1–Ev4 (i.e., except those related to the heliport)  
 858 are also present in videos 9 and 10. Since these videos are  
 859 set on a crossroad, events Ev3 and Ev4 more frequently are  
 860 related to vehicles. In particular, Ev4 can refer to vehicle-  
 861 person accidents, as well as only vehicle accidents. The system  
 862 performance on discovering and identifying the critical situa-  
 863 tions has been evaluated by the accuracy measure, defined as  
 864 the number of discovered (critical) events, detected per scene  
 865 by our framework, with respect to the total number of events,  
 866 defined by the annotated ground truth. More formally, given a  
 867 scene  $s$ , the accuracy  $A_e$  on an event  $e$  occurring in the scene  
 868  $s$ , is the number of detected events  $S_e$  that match the total  
 869 number of defined events  $E_e$

$$870 \quad A_e = \frac{S_e \cap E_e}{E_e}. \quad (1)$$

871 Table II shows test results on the situation recognition by con-  
 872 sidering the scenes from the videos, presented in the previous  
 873 section. The accuracy  $A_e$  is evaluated for all the listed critical  
 874 events. Events marked with “/” symbol do not occur or are  
 875 not allowed for that specific video.

876 The experimental results are promising: critical events  
 877 occurring in the videos are often detected mainly because of  
 878 an efficient reasoning process and the high-quality data pro-  
 879 vided by the geometric tools, as well as the used map-based  
 880 tools. These results are likely due to the nature of videos that  
 881 are limited to two types of moving objects: 1) persons and  
 882 2) vehicles. Moreover, the reasoning model has been “tailored”  
 883 to detect alerting situations, when people and vehicles come  
 884 into play. Undiscovered critical situations may depend on var-  
 885 ious factors. First, an object that is not correctly retrieved for  
 886 many frames in the scene can affect the detection of an alert-  
 887 ing event. Then, a critical event is mainly related to object  
 888 GPS coordinates, thus, if they are not exact, some positional  
 889 relations (e.g., inTheAreaOf, inFrontOf, etc.) could be missing  
 890 and consequently, a critical event could not be detected.

891 The accuracy for video 1 is satisfactory in all the scenes:  
 892 in scene 1 [Fig. 5(a)], the two students coming near the route  
 893 are correctly detected; in scene 2 [Fig. 5(b)] the system recog-  
 894 nizes the two students and the third guy walking, respectively,  
 895 near and around the route. In scene 3 [Fig. 5(c)], one student  
 896 crossing the road twice is detected: the first time when a car  
 897 is coming and the second time when the road is empty; the  
 898 system also correctly detects the students while standing near  
 899 the road. In scene 4, the system detects the students com-  
 900 ing near the route, but a student walking on the sidewalk is  
 901 detected as crossing the road, due to the low location accuracy  
 902 of GPS Google Maps.

TABLE II  
 SCENE UNDERSTANDING TEST RESULTS

Video	Scenes	Accuracy ( $A_e$ )					
		Ev1	Ev2	Ev3	Ev4	Ev5	Ev6
1	1	1.0	1.0	/	/	/	/
	2	1.0	1.0	/	/	/	/
	3	/	1.0	1.0	1.0	/	/
	4	1.0	0.67	/	/	/	/
2	1	1.0	1.0	/	/	1.0	1.0
	2	/	1.0	/	/	1.0	1.0
	3	/	/	/	/	1.0	1.0
	4	/	/	/	/	1.0	1.0
3	1	1.0	1.0	/	/	/	/
	2	1.0	1.0	/	/	/	/
	3	1.0	1.0	/	/	/	/
	4	1.0	1.0	1.0	1.0	/	/
4	1	1.0	/	/	/	/	/
	2	1.0	/	/	1.0	/	/
	3	1.0	/	1.0	1.0	/	/
	4	1.0	0.82	1.0	/	/	/
6	1	/	/	/	1.0	/	/
	2	/	/	/	/	/	/
	3	/	/	/	/	/	/
7	1	/	/	/	1.0	/	/
	2	/	/	/	1.0	/	/
	3	/	/	/	/	/	/
8	1	/	/	/	1.0	/	/
	2	/	/	/	/	/	/
	3	/	/	/	1.0	/	/
9	1	1.0	1.0	/	1.0	/	/
	2	0.96	1.0	/	1.0	/	/
	3	1.0	1.0	/	1.0	/	/
10	1	0.92	0.87	1.0	1.0	/	/
	2	1.0	1.0	1.0	1.0	/	/
	3	1.0	1.0	1.0	1.0	/	/

Experimental results in video 2 confirm the good 903 performance of the system in event identification. In scene 1 904 [Fig. 6(a)], a student coming toward the center of the heliport 905 and other two persons standing near the heliport are correctly 906 detected. In scenes 2 and 3 [Fig. 6(b) and (c)], the two stu- 907 dents walk toward the center of the heliport and then stop for 908 a while. The system still correctly recognizes the two events: 909 walking on the border and on the center of the landing area. 910 Both events are regarded as very critical because the students 911 are inside the heliport area. Finally, in scene 4 [Fig. 6(d)], 912 the system correctly identifies the other student coming to the 913 border of the heliport. 914

Similarly to video 1, video 3 shows, in the first three scenes, 915 a person walking on the lawn away from the route and the two 916 students going toward the route; the fourth scene presents one 917 of the two students crossing with and without a car coming. 918 Also in these cases, the system successfully detects all the 919 critical situations. 920

Video 4 also shows initially only a student walking on 921 the lawn, then it shows another student crossing the route in 922 scenes 2 and 3. In the final scene, there are the two students 923 stopping near the road. Situations are correctly recognized 924 except some Ev2 events in the fourth scene due to the low 925 location accuracy of GPS Google Maps. 926

For videos 5–8, the system detects very few critical situa- 927 tions; these videos present just a single type of object: Person 928 (video 5) or Vehicle (videos 6–8), moving in safe conditions 929

in a flat scenario with a very few elements. Video 5 indeed shows two people walking on a deserted lawn and doing nothing else. In fact, no event is detected (e.g., no result is presented in the table). Video 6 instead, shows some cars running around a roundabout: the only alerting situation occurs between two cars in the first scene, where a presumed collision (Ev4) was about to happen, and the system correctly recognizes it. Videos 7 and 8 present route events similar to those seen in video 6, with a consequent similar accuracy value.

Videos 9 and 10 are rich of events happening in the same time after a car crashing, with a car stopped in the middle of the crossroad. The results on video 9 are very good: in scene 1, the system correctly recognizes, as a critical situation, the occurrence of events Ev2 and Ev4, describing a person leaving the car and crossing while more cars are coming. Some other people walking near the road and crossing are also successfully detected (Ev1). In scene 2, the situation describing some cars going toward the stopped car is successfully recognized as critical, as well as some people walking near the road when cars are coming. In scene 3, all the events Ev4 are detected: they mainly refer to cars that invade the side of the oncoming traffic roadway. Finally, in video 10, scene 1, there are many different co-occurring events: people close to the road and crossing toward oncoming vehicles (Ev1, Ev2, and Ev4), cars and trucks that quickly drive into the oncoming traffic (Ev4). The system identifies the most of the occurring events, only few Ev1 and Ev2 events are lost.

## VII. CONCLUSION

This paper introduces a novel approach to enhance the video tracking task, by adding semantic information to the tracks appearing in the video frames, in order to provide a high-level comprehensive description of the scenes recorded by the mobile camera video. The final goal is to get a more accurate object identification and labeling and, at the same time, to get critical situations and scene understanding based on the identified relations. The proposed framework exploits the semantic Web technologies to extend the assertional knowledge extracted by the video analysis. The reasoning model has been designed to infer new assertions that support the understanding of the dynamics in the scene evolution.

The experimentation in real applications shows interesting results on videos recorded with a camera-on-board drone. The add-on is the use of mobile camera, compared with the most works in literature that are based on fixed camera.

Future extensions will focus on improving the reliability of the system in object recognition and the dynamic estimation of the dangers affecting a specific scenario. The aim is providing an “alarm” measure, that allows identifying potential risks during the development of the scene: each identified scenario will be comprehensively described by a danger value which represents the alarming level detected in the scene.

## REFERENCES

- [1] S. Saripalli, J. F. Montgomery, and G. S. Sukhatme, “Visually guided landing of an unmanned aerial vehicle,” *IEEE Trans. Robot. Autom.*, vol. 19, no. 3, pp. 371–380, Jun. 2003.
- [2] A. Symington, S. Waharte, S. Julier, and N. Trigoni, “Probabilistic target detection by camera-equipped UAVs,” in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, Anchorage, AK, USA, May 2010, pp. 4076–4081.
- [3] M. Ugliano, L. Bianchi, A. Bottino, and W. Allasia, “Automatically detecting changes and anomalies in unmanned aerial vehicle images,” in *Proc. IEEE 1st Int. Forum Res. Technol. Soc. Ind. Leveraging Better Tomorrow (RTS)*, Turin, Italy, Sep. 2015, pp. 484–489.
- [4] M. Siam, R. ElSayed, and M. ElHelw, “On-board multiple target detection and tracking on camera-equipped aerial vehicles,” in *Proc. IEEE Int. Conf. Robot. Biomimetics (ROBIO)*, Guangzhou, China, Dec. 2012, pp. 2399–2405.
- [5] R. Cui, C. Yang, Y. Li, and S. Sharma, “Adaptive neural network control of AUVs with control input nonlinearities using reinforcement learning,” *IEEE Trans. Syst., Man, Cybern., Syst.*, vol. 47, no. 6, pp. 1019–1029, Jun. 2017.
- [6] K. Dorling, J. Heinrichs, G. G. Messier, and S. Magierowski, “Vehicle routing problems for drone delivery,” *IEEE Trans. Syst., Man, Cybern., Syst.*, vol. 47, no. 1, pp. 70–85, Jan. 2017.
- [7] S. Kwak, M. Cho, I. Laptev, J. Ponce, and C. Schmid, “Unsupervised object discovery and tracking in video collections,” in *Proc. IEEE Int. Conf. Comput. Vis.*, Santiago, Chile, 2015, pp. 3173–3181.
- [8] F. D. Smedt, D. Hulens, and T. Goedemé, “On-board real-time tracking of pedestrians on a UAV,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. Workshops (CVPRW)*, Boston, MA, USA, Jun. 2015, pp. 1–8.
- [9] S. Minaeian, J. Liu, and Y.-J. Son, “Vision-based target detection and localization via a team of cooperative UAV and UGVs,” *IEEE Trans. Syst., Man, Cybern., Syst.*, vol. 46, no. 7, pp. 1005–1016, Jul. 2016.
- [10] C.-Y. Hsu, L.-W. Kang, and H.-Y. M. Liao, “Cross-camera vehicle tracking via affine invariant object matching for video forensics applications,” in *Proc. IEEE Int. Conf. Multimedia Expo (ICME)*, San Jose, CA, USA, Jul. 2013, pp. 1–6.
- [11] J. T. Lee, C.-C. Chen, and J. K. Aggarwal, “Recognizing human-vehicle interactions from aerial video without training,” in *Proc. IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit. Workshops (CVPRW)*, Colorado Springs, CO, USA, Jun. 2011, pp. 53–60.
- [12] J. Wang, J. Qian, and R. Ma, “Urban road information extraction from high resolution remotely sensed image based on semantic model,” in *Proc. 21st Int. Conf. Geoinformatics*, Kaifeng, China, Jun. 2013, pp. 1–5.
- [13] H. Zhou, H. Kong, L. Wei, D. Creighton, and S. Nahavandi, “Efficient road detection and tracking for unmanned aerial vehicle,” *IEEE Trans. Intell. Transp. Syst.*, vol. 16, no. 1, pp. 297–309, Feb. 2015.
- [14] V. Bruni and D. Vitulano, “An improvement of kernel-based object tracking based on human perception,” *IEEE Trans. Syst., Man, Cybern., Syst.*, vol. 44, no. 11, pp. 1474–1485, Nov. 2014.
- [15] T. Chen and S. Lu, “Object-level motion detection from moving cameras,” *IEEE Trans. Circuits Syst. Video Technol.*, to be published, doi: 10.1109/TCSVT.2016.2587387.
- [16] P. Huang, J. Cai, Z. Meng, Z. Hu, and D. Wang, “Novel method of monocular real-time feature point tracking for tethered space robots,” *J. Aerosp. Eng.*, vol. 27, no. 6, 2013, Art. no. 04014039.
- [17] L. Chen, P. Huang, J. Cai, Z. Meng, and Z. Liu, “A non-cooperative target grasping position prediction model for tethered space robot,” *Aerospace Sci. Technol.*, vol. 58, pp. 571–581, Nov. 2016.
- [18] J. Cai, P. Huang, B. Zhang, and D. Wang, “A TSR visual servoing system based on a novel dynamic template matching method,” *Sensors*, vol. 15, no. 12, pp. 32152–32167, 2015.
- [19] M. Piccardi, “Background subtraction techniques: A review,” in *Proc. IEEE Int. Conf. Syst. Man Cybern.*, vol. 4. The Hague, The Netherlands, Oct. 2004, pp. 3099–3104.
- [20] R. Brehar, C. Fortuna, S. Bota, D. Mladenic, and S. Nedevschi, “Spatio-temporal reasoning for traffic scene understanding,” in *Proc. IEEE Int. Conf. Intell. Comput. Commun. Process. (ICCP)*, Cluj-Napoca, Romania, Aug. 2011, pp. 377–384.
- [21] M.-C. Chuang, J.-N. Hwang, J.-H. Ye, S.-C. Huang, and K. Williams, “Underwater fish tracking for moving cameras based on deformable multiple kernels,” *IEEE Trans. Syst., Man, Cybern., Syst.*, vol. 47, no. 9, pp. 2467–2477, Sep. 2017.
- [22] M. Eldib *et al.*, “A low resolution multi-camera system for person tracking,” in *Proc. IEEE Int. Conf. Image Process. (ICIP)*, Paris, France, Oct. 2014, pp. 378–382.
- [23] X. Li and H. Lu, “Object tracking based on local learning,” in *Proc. 19th IEEE Int. Conf. Image Process.*, Orlando, FL, USA, Sep. 2012, pp. 413–416.

- 1061 [24] W. Jiang, C. Xiao, H. Jin, S. Zhu, and Z. Lu, "Vehicle tracking with  
1062 non-overlapping views for multi-camera surveillance system," in *Proc.  
1063 IEEE 10th Int. Conf. High Perform. Comput. Commun. IEEE Int. Conf.  
1064 Embedded Ubiquitous Comput.*, Nov. 2013, pp. 1213–1220.  
1065 [25] Z. Zhang and F. Cohen, "3D pedestrian tracking based on overhead  
1066 cameras," in *Proc. 7th Int. Conf. Distrib. Smart Cameras (ICDSC)*,  
1067 Palm Springs, CA, USA, Oct. 2013, pp. 1–6.  
1068 [26] Y. Yuan, L. Mou, and X. Lu, "Scene recognition by manifold regular-  
1069 ized deep learning architecture," *IEEE Trans. Neural Netw. Learn. Syst.*,  
1070 vol. 26, no. 10, pp. 2222–2233, Oct. 2015.  
1071 [27] C. Shen, N. Yu, W. Li, and W. Zhou, "Moving object detection  
1072 of dynamic scenes using spatio-temporal context and background  
1073 modeling," in *Proc. 6th Int. Conf. Wireless Commun. Signal Process.  
1074 (WCSP)*, Hefei, China, Oct. 2014, pp. 1–4.  
1075 [28] C. F. Crispim-Junior *et al.*, "Semantic event fusion of different visual  
1076 modality concepts for activity recognition," *IEEE Trans. Pattern Anal.  
1077 Mach. Intell.*, vol. 38, no. 8, pp. 1598–1611, Aug. 2016.  
1078 [29] C. Kothari, J. Qualls, and D. Russomanno, "An ontology-based data  
1079 fusion framework for profiling sensors," in *Proc. IEEE Int. Conf.  
1080 Electro/Inf. Technol. (EIT)*, Indianapolis, IN, USA, May 2012, pp. 1–6.  
1081 [30] J. Gomez-Romero, M. A. Patricio, J. Garcia, and J. M. Molina,  
1082 "Context-based reasoning using ontologies to adapt visual tracking in  
1083 surveillance," in *Proc. 6th IEEE Int. Conf. Adv. Video Signal Based  
1084 Surveillance (AVSS)*, Genoa, Italy, Sep. 2009, pp. 226–231.  
1085 [31] M. Andersson *et al.*, "Activity recognition and localization on a truck  
1086 parking lot," in *Proc. 10th IEEE Int. Conf. Adv. Video Signal Based  
1087 Surveillance (AVSS)*, Krakow, Poland, Aug. 2013, pp. 263–269.  
1088 [32] X. Zhang *et al.*, "Human pose estimation and tracking via parsing a tree  
1089 structure based human model," *IEEE Trans. Syst., Man, Cybern., Syst.*,  
1090 vol. 44, no. 5, pp. 580–592, May 2014.  
1091 [33] J. Garcia, A. Gardel, I. Bravo, J. L. Lázaro, and M. Martinez, "Tracking  
1092 people motion based on extended condensation algorithm," *IEEE Trans.  
1093 Syst., Man, Cybern., Syst.*, vol. 43, no. 3, pp. 606–618, May 2013.  
1094 [34] S. Dasiopoulou, V. Mezaris, I. Kompatsiaris, V.-K. Papastathis, and  
1095 M. G. Strintzis, "Knowledge-assisted semantic video object detec-  
1096 tion," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 15, no. 10,  
1097 pp. 1210–1224, Oct. 2005.  
1098 [35] H. He and B. Upcroft, "Nonparametric semantic segmentation for 3D  
1099 street scenes," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, Tokyo,  
1100 Japan, Nov. 2013, pp. 3697–3703.  
1101 [36] Y. Li, Y. Guo, Y. Kao, and R. He, "Image piece learning for weakly  
1102 supervised semantic segmentation," *IEEE Trans. Syst., Man, Cybern.,  
1103 Syst.*, vol. 47, no. 4, pp. 648–659, Apr. 2017.  
1104 [37] D. Cavaliere, S. Senatore, M. Vento, and V. Loia, "Towards semantic  
1105 context-aware drones for aerial scenes understanding," in *Proc. 13th  
1106 IEEE Int. Conf. Adv. Video Signal Based Surveillance (AVSS)*, Colorado  
1107 Springs, CO, USA, Aug. 2016, pp. 115–121.  
1108 [38] A. Greco, A. Saggese, G. Tozzi, and M. Vento, "Detection and tracking  
1109 of moving objects by flying cameras using an adaptive test and set  
1110 approach," in *Proc. 22nd Comput. Vis. Win. Workshop (CVWW)*, 2017,  
1111 pp. 1–6.  
1112 [39] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and  
1113 A. Zisserman, "The Pascal visual object classes (voc) challenge," *Int. J.  
1114 Comput. Vis.*, vol. 88, no. 2, pp. 303–338, 2010.



**Vincenzo Loia** (SM'08) received the B.S. degree 1127 from the University of Salerno, Fisciano, Italy, in 1128 1985, and the M.S. and Ph.D. degrees from the 1129 University of Paris VI, Paris, France, in 1987 and 1130 1989, respectively, all in computer science. 1131

Since 1989, he has been a Faculty Member with 1132 the University of Salerno, where he teaches safe 1133 systems, situational awareness, IT project and ser- 1134 vice management and is currently the Chair and a 1135 Professor of computer science with the Department 1136 of Management and Innovation Systems. His current 1137 research interests include soft computing, agent technology for technologically 1138 complex environments Web intelligence, and situational awareness. He was 1139 a principal investigator in industrial research and development projects and 1140 in academic research projects. He has authored over 400 original research 1141 papers in international journals, book chapters, and international conference 1142 proceedings. 1143

Dr. Loia is the Coeditor-in-Chief of *Soft Computing* (Springer) and 1144 the Editor-in-Chief of *Ambient Intelligence and Humanized Computing* 1145 (Springer). He is an Associate Editor of various journals, including the 1146 IEEE TRANSACTIONS ON SYSTEM, MAN AND CYBERNETICS: SYSTEMS, 1147 the IEEE TRANSACTIONS ON FUZZY SYSTEMS, the IEEE TRANSACTIONS 1148 ON INDUSTRIAL INFORMATICS, the IEEE TRANSACTIONS ON COGNITIVE 1149 AND DEVELOPMENTAL SYSTEMS. He holds several roles in the IEEE Society 1150 in particular for Computational Intelligence Society (the Chair of the Emergent 1151 Technologies Technical Committee, IEEE CIS European Representative, and 1152 the Vice-Chair of the Intelligent Systems Applications Technical Committee). 1153



**Alessia Saggese** (M'14) received the Ph.D. degree 1154 in computer engineering from the University of 1155 Salerno, Fisciano, Italy, and the University of Caen 1156 Normandy, Caen, France, in 2014. 1157

She is currently an Assistant Professor with 1158 the University of Salerno. Her current research 1159 interests include basic methodologies and appli- 1160 cations in computer vision and pattern recogni- 1161 tion. 1162

Dr. Saggese has been a member of the 1163 International Association for Pattern Recognition 1164 15 on Graph-Based Representations in Pattern 1165 1166

Technical Committee  
Recognition since 2012.



**Sabrina Senatore** (M'06) received the M.S. and 1167 Ph.D. degrees in computer science from the 1168 University of Salerno, Fisciano, Italy, in 1999 and 1169 2004, respectively. 1170

Since 2005, she has been a Faculty Member with 1171 the University of Salerno, where she is an Associate 1172 Professor of computer science. Her current research 1173 interests include the development and application 1174 of intelligent systems based on the combination 1175 of techniques from soft computing, computational 1176 intelligence, text mining, Web information retrieval, 1177 intelligent systems, and machine learning, areas in which she 1178 has published numerous papers. 1179

Dr. Senatore is an Editorial Board Member of the *Applied Intelligence* 1180 Journal and the *International Journal of Computational Intelligence Theory* 1181 and Practice. 1182



**Mario Vento** (M'14) received the Ph.D. degree in 1183 computer engineering from the University of Napoli 1184 Federico II, Naples, Italy, in 1989. 1185

He is currently a Full Professor of computer sci- 1186 ence and artificial intelligence with the University of 1187 Salerno, Fisciano, Italy, where he is the Coordinator 1188 of the Artificial Vision Laboratory and the Dean 1189 of the Department of Information and Electrical 1190 Engineering and Applied Mathematics. His current 1191 research interests include real-time video analysis 1192 and interpretation for video surveillance applica- 1193 tions, classification techniques, either statistical, syntactic, and structural, exact 1194 and inexact graph matching, multiexpert classification, and learning method- 1195 ologies for structural descriptions. 1196

Prof. Vento served as the Chairman of the International Association Pattern 1197 Recognition (IAPR) Technical Committee 15 on Graph-Based Representation 1198 in Pattern Recognition from 2002 to 2006. He is a Fellow Scientist of the 1199 IAPR. 1200



**Danilo Cavaliere** received the master's degree in 1115 computer science from the University of Salerno, 1116 Fisciano, Italy, in 2014, where he is currently pur- 1117 suing the Ph.D. degree. 1118

Since 2015, he has been a Pre-Doctoral 1119 Fellow with the Computer Science Department 1120 and the Department of Information and Electrical 1121 Engineering and Applied Mathematics, University 1122 of Salerno. His current research interests include 1123 artificial and computational intelligence, machine 1124 learning, data mining, and knowledge discovery, 1125 areas in which he has published some papers. 1126