

Semantic-driven modelling of context and entity of interest profiles for maritime situation awareness

Elena CAMOSSÌ, Francesca DE ROSA

Science and Technology Organization, Centre for Maritime Research and Experimentation, La Spezia, Italy^{1,2}

Abstract One of the inherent aspects of Situation Awareness (SA) is the correct interpretation of the perceived situational picture, to enable future projection and support decision making [1]. In surveillance, security operators must be able to focus on the most important events in the picture, and to evaluate the threat risk, which is assessed in relation to the event context. The notion of context has long been investigated in pervasive computing, mostly for Internet of Things and computer SA. In this position paper, we propose a formalisation of SA context, applicable to events interpretation for security and safety threat assessment. We exemplify it for Maritime SA (MSA), to contextualise maritime events, facts and anomalies, and assess the risk associated to maritime threats. The formalisation relies on entity profiles, which represent the relevant historical knowledge on the entities of interest for security, i.e., vessels, areas, information sources. Profiles are built over time, updating and elaborating the observations generated by the maritime sensor network, an approach suitable to Linked Data and Knowledge Graphs.

Keywords. Context, Events, Occurents, Situational Awareness, Heterogeneous Sensor Network, Surveillance, Security, Anomalies, Ontology, Knowledge Graph

1. Introduction

Context modelling and reasoning have received significant attention in the area of pervasive and ubiquitous computing, with a specific focus on situation awareness and context awareness of intelligent/smart agents and Internet of Things (IoT). Situational Awareness (SA) represents the human mental model of a situation and is obtained through Situational Assessment, which is the process that seamlessly condenses the acquisition of information from the real world, its interpretation to understand the ongoing situation and the projection in the future to support decision making [1]. Although proposed for human SA, this concept in general appears to be applicable also to systems [2].

¹Copyright© 2020 for this paper by its authors. Authors affiliated with the NATO STO CMRE are employees of the North Atlantic Treaty Organization (NATO) who have prepared this work as part of their official duties. Accordingly, NATO retains non-exclusive royalty-free rights of total or partial reproduction, dissemination and preparation of derivative works based on this work, for NATO official purposes. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

²The authors acknowledge the support of the EU H2020 INFORE project (GA No. 825070).

In this paper, we address the formalisation of context in surveillance applications that can be integrated in information fusion systems in support of event interpretation for situation comprehension and threat assessment. Moreover, we will describe its application to the maritime domain.

Context and situation are intertwined concepts and the artificial intelligence community debates on their definition since the late 80s [3, 4]. Several distinct formalisation approaches are presented in the literature (e.g., [4–8]). Two perspectives relevant to this work are proposed in [9] and [10]. The first one addresses human SA and investigates context for situation assessment in decision aid systems. It distinguishes between internal context (i.e., agents cognitive state) and the external context, which addresses the situation and the understanding of the environment [10]. Events in [9] are situational elements, which can be combined in situations. In [10], context (application) spaces are multidimensional spaces whose dimensions are variables observed by sensors, and a context state is a multidimensional point in the application space.

With respect to existing contribution in the maritime domain (e.g. [11]) and others, we focus on the representation of external context, while we do not discuss the internal context of the agent (i.e., the surveillance operator). The only internal context factor considered is the mission, as it sets the goals of the agent. The proposed formalisation is based on the profiles of observed entities, which encode the historical knowledge on the entities of interests. Profiles are constructed incrementally, aggregating and analysing the observations acquired through (possibly heterogeneous) sensor networks. These profiles provide the context to help the agent interpret events and assess potential threats. In fact, similarly to [10], the context supports the agent in distinguishing among different situations, all feasible on the basis of the detected events. Differently than in [10], in our proposal threats correspond to situations (and have corresponding situational spaces), but the context space is composed by all information produced by sensors, including analytics (e.g., the analytics that detect and forecast events).

This formalisation suites loosely structured models like Knowledge Graphs, where new information on an entity can be simply added to its graph and the associated profile information can be updated accordingly. To demonstrate the feasibility of the proposed approach, we exemplify our proposal with an ontology-based design for maritime event context and entity profiles.

The remainder of this paper is organised as follows. Section 2 presents the underlying idea of information profiles and context for threats assessment, while Section 3 illustrates and exemplifies an ontology-driven design of SA context. . Finally, Section 4 concludes the paper, discussing future research directions.

2. From entity profiles to context for threat assessment

The idea of context we propose in this paper raises from the observation that profile information, a concept that is largely applied in security, frames a reasoned representation of relevant historical knowledge that expresses the threat risk associated to scenario entities (e.g., the risk of collision in an area), and which is used to assess the potential risk to security (or safety) of a situation.

The SA scenario addressed in this paper comprises an agent who is monitoring an area for security (or safety) purposes. The agent is supported by a fusion system con-

ected to a sensor network composed of a variety of surveillance systems. The agent receives from the fusion system notifications on a series of events, which are detected and forecasted as potential indicators of threatening situations (e.g., trafficking, smuggling, collision). For instance, vessels stopping in or approaching interdicted areas, abruptly changing direction or anchoring close to other vessels outside ports can be considered out of the common situations, and would be signalled by the system. Given the set of alarms, the agent assesses the situation to decide if an intervention is needed.

The agent needs additional contextual information to perform the situational assessment.

Specifically, the agent assesses the reliability of the surveillance network in the area (e.g., is the sensor coverage good? are interference or malfunctions frequent?), the risk profile of the area (e.g., are there smuggling reports? does it include critical infrastructures?), and the history of the vessels raising the alarms (e.g., where did those vessels ship before? where do they declare to go?). The profile information of the entities (sensors, area, vessels) associated to the alarms, once combined, form the events context and are used by the agent to assess the likelihood of maritime threats.

Entity profiles and event context may be constructed building on basic information items. In SA, these are obtained by fusing and elaborating sensor observations. Observations are collected over time, and knowledge is derived with analytics, e.g., producing statistics on threats risk. In a SA scenario as the one illustrated above, entity profiles condense historical knowledge on entities, sensor quality, patterns-of-life (i.e., expected behaviours like the usual traffic in an area), statistics on threats and events. When an event is detected, it is sufficient to exploit the profile of the associated entities to get the event's context and assess the threat risk.

3. Representing entity profiles and events context in a SA ontology

In this section, we illustrate how the information model of a typical information fusion system can be extended to support entity profiles and event context for SA. The starting information model is the MSA Heterogeneous Sensor Network (MSA-HSN) ontology [12], which has been defined to semantically annotate the information generated by the maritime use case of the H2020 INFORE project³. The MSA-HSN includes information components for sensors, observations, measures, features of interest, events, quality of information and sources, and information provenance. It extends and adapts existing information models for sensors and observations, in particular the Semantic Sensor Network /Sensors, Observations, Samples, Actuators ontology (SSN/SOSA) for sensors and sensor observations [13], and leverages the specification of maritime data and events of the Common Information Sharing Environment (CISE) data exchange model [14].

The top-level concepts of MSA-HSN are represented in Figure 1. In the MSA-HSN, all information sources are modelled as sensors (Figure 1, top-left). Sensors observe properties of entities (namely, features of interest, like vessels) and produce observations by executing procedures (e.g. sensor plans, algorithms). Observations are composed by qualitative and quantitative measures (Figure 1, top-right), have spatio-temporal characteristics as well as associated quality evaluations (e.g., confidence). Similarly, sensors

³Interactive Extreme-Scale Analytics and Forecasting (INFORE) project website: www.info-re-project.eu




Figure 1. The Maritime Situational Awareness Heterogeneous Sensor Network Ontology, extended to support entity profiles and event context. Boxes represent top-level ontology concepts, including sensors and observations (violet background), observation values (light blue), and maritime events (orange). Arrows represent object properties linking concepts. Empty boxes are concepts from existing ontologies (e.g., [13, 14]).

quality (e.g., reliability) may be represented as an observation. Observations may also derive from existing observations. Sensors and originating observations are used to model the observation provenance.

The fusion system produces event notifications, using a variety of approaches like machine learning, signal processing, statistical analysis and rule-based approaches. Events (Figure 1, bottom, concept *Event*) are situational elements [9] that exist for a limited period of time and are bounded in space, like processes, phenomena and activities. They may be simple or complex, like activities (e.g., fishing).

At each instant in time, a set of events is detected and forecasted, producing event notifications. An event is temporally bounded, occurs in a location (a geo-referenced spatial value, which can be associated to an area of interest like a port), and applies to a feature of interest, e.g. a vessel. As illustrated in Figure 1, an event notification is a particular type of observation (see the modelling of *Event* and *MSAObservation*⁴). As such, it is generated by a sensor (e.g., an anomaly detection or complex event processing software). A set of observations (or signals) generated by other sensors is analysed to produce an event. These may include raw sensor information and other alarms. The event can also be associated to an agent, and to other related events.

⁴In the MSA-HSN and in Figure 1, instances of *Event* correspond to event notifications.

3.1. Event context and threat assessment

Each event detected by the fusion system is contextualised by a combination of information. To start with, the agent looks for some descriptive information on the event's object (e.g., information available from vessel registers). Then, they take into consideration the object history (e.g., the past movements of the vessel) and its profile. This includes the object patterns-of-life, modelling the expected object behaviour (e.g., the vessel most visited ports, most frequent routes) and the object risk indicators (e.g., the vessel incident reports). The profiles of all the areas of interest overlapping the event location, and the profiles of the sensors that produced the raw signals can also be taken into account.

The bottom, dark orange concepts in Figure 1 shown how the MSA-HSN ontology has been extended with profile classes to formalise threat, event, sensor, location, and feature of interest profiles. The specific profile representations may differ, as described in this section, but a profile may encompass other profile types, with specific characteristics. For instance, the profile of an area of interest *LocationProfile* includes the event and threat profiles for the area (i.e., statistical information on the occurrence of threats and events in the area, which can be linked to incident reports), and sensor profiles (i.e., the performance evaluation of the sensors that observe - and observed - the area). The profiles of areas and features of interest include additional information, like patterns-of-life to characterise typical entity behaviours (e.g., expected traffic), information on infrastructures and human activities in the area, etc.

Given this extended model, any new observation received through the sensor network is added to the knowledge graph, and the associated information profiles are updated. In the case of a vessel, the associated voyage history and pattern-of-life are updated to take into account the recent movements; if the vessel is involved in some incident, its risk profile is amended. Sensors' quality is already formalised in the ontology, and the relevant indices (e.g. coverage, reliability) are linked to the sensor profiles. Similarly, the profiles of areas of interest are updated with the last available information on maritime traffic (aggregated from vessel movement observations) and sensors. Threat and event profiles are updated with the information available on vessels and areas of interest.

The class *Context* is modelled in the ontology on top of profiles and historical information. Whenever an event is added to the knowledge graph, an event context instance is created and associated to the event, embedding the relevant profile information, readily available to support the end user of the system in assessing the situation.

The threat assessment is performed as follows. In a preliminary assessment, a set of potential threats is selected, by comparing the temporally valid events with the behavioural characteristics of the threats, which are expressed as event patterns (e.g., the vessel stopped in an interdicted area, disappeared from the sensor network, then reappeared after few hours outside the interdicted area).

Then, the set of potential threats is assessed against the context of the valid events, by comparing the event context with the threat profiles. A threat profile expresses the likelihood of the threat, per areas of interest and feature of interest categories (e.g., vessel types, vessel flags). Together with the event context, the threat profile allows to assess the likelihood of the threat for the particular situation. Finally, if some a-priori intent information is available and is verified, the assessment could shift the belief towards a security threats instead of a safety one.

4. Conclusions

This paper proposes a semantic modelling of context for SA, applicable to security and safety threat assessment. The formalisation relies on profile modelling of entities of interest, to encapsulate historical knowledge useful to evaluate the threat risk and likelihood. The application of the proposed formalisation is described, and an ontology driven modelling is illustrated. Future work includes the evaluation of the proposed formalisation approach against the real world maritime security use case defined in INFORE, using the ontology model presented in the paper.

References