

Predicting Terroristic Attacks in Urban Environments: An Internet-of-Things Approach

Stavros Petris (spet@ait.gr), Christos Georgoulis (cgeo@ait.edu.gr) and John Soldatos (jsol@ait.edu.gr)

*Athens Information Technology, 0,8km Markopoulo Ave., P.O. Box 68, GR-19002
Peania, Greece*

Ilaria Giordani (giordani@milanoricerche.it), Raul Sormani
(giordani@milanoricerche.it) and Divna Djordjevic (djordjevic@milanoricerche.it)
Consorzio Milano Ricerche, Via Cozzi 53, 20125 Milano, Italy

Abstract

In the recent years we have witnessed a number of important terroristic incidents, in major cities all around the world (e.g., 911 in New York, 11-M in Madrid, 7/7 in London). These incidents have revealed the vulnerabilities of urban environments, against terroristic plans and have created significant pressure towards devising novel tools and techniques for timely predicting the intentions and plans of terrorists. In this paper, we introduce a blueprint Internet-of-Things architecture for predicting terroristic attacks. The architecture allows Law enforcement agencies to exploit multiple data sources, (including SIGINT, OSINT and HUMINT) towards acquiring information associated with terroristic action, while at the same time providing powerful reasoning capabilities towards transforming raw events into meaningful alerts. We also illustrate the implementation of a terroristic prediction system based on this architecture, along with its use in the scope of a validating scenario.

Keywords: *Urban Environment, Sensors, Internet-of-Things, Terrorist Indicators, Predictive Reasoning*

1. Introduction

During the last fifteen years the world has witnessed major terroristic attacks and security incidents in some of the most important cities of the world. Prominent examples include the 911 attack in NYC's (New York City) world trade center, the train bombings in Madrid (also known in Spain as 11-M), as well as the suicide attacks in the London underground (also referred to as 7/7). These incidents have exposed the vulnerabilities of the urban environment against actions of terrorists, which mainly stem from its diversity, heterogeneity and complexity. Indeed, the presence of civilians, the availability of many and diverse physical infrastructures, as well as the complex social, cultural, and governmental interactions that support urban life, tend to provide room for terrorists to plan and commit their attacks, while at the same time security and Law Enforcement Agencies (LEAs) have very hard time in reacting to these attacks. In this context, LEAs are in need of novel tools and techniques that could essentially help them in identifying terroristic plans as well as in anticipating terroristic actions.

Key to the early identification of terroristic actions, but also to the possible prediction of terroristic attacks, is the collection and exploitation of information from the urban environment, notably, information that could be associated with either the preparatory or the

operational phases of terroristic attacks. The collection of this information is typically based on a variety of sources including: (A) Signal Intelligence (SIGINT) sources, *i.e.*, information/intelligence derived from the interception and combination of signals (*e.g.*, stemming from cellular phones, fax, and radio), (B) Imaginary Intelligence (IMINT) sources, which refers to intelligence derived from satellites, cameras and aerial photography (including Unmanned Aerial Vehicles (UAVs)), (C) Human Intelligence (HUMINT)[8], *i.e.*, intelligence based on information collected and provided by human sources including both obvious (overt) and secret (clandestine) sources and (D) Open Source Intelligence (OSINT) [14], which refers to intelligence that is based on unclassified public sources (such as books, technical manuals, asset websites, but also emerging social media (blogs, social networks, *etc.*)). While human sources (such as patrolling policemen and officers) can provide accurate information about unusual activities and events, sensors and computer-based sources can be used to obtain information beyond the capacity of the human resources of the LEAs. Indeed, sensors can be used to monitor and obtain information from multiple areas within a city, without the need for patrolling these areas. Furthermore, open sources (such as social networks) provide abundant information that can complete information derived from other sources. The collection of information about possible terroristic activities and events is not sufficient to lead to the prediction of terroristic attacks. This is because information and events derived from the above sources may include events and information unrelated to terrorist attacks. To this end, there is a clear need for analyzing the collected information towards identifying events and behaviors that are highly likely to be linked with terroristic activities.

Recent advances in ICT and more specifically in multi-sensor systems and Big Data analytics enable the development of systems that can collect and process information from a wide variety of sources, including structured and unstructured data, but also real-time and non-real time data. Therefore, such a system can serve as a basis for collecting information from multiple heterogeneous sources (including sensors and information databases) and accordingly executing analytics algorithms that could extract and assess potential terroristic activities. Closely related to multi-sensor systems is the internet-of-things paradigm [17], which enables the orchestration and coordination of a large number of physical and virtual Internet-Connected-Objects (ICO) towards human-centric services in a variety of sectors including logistics, trade, industry, smart cities and ambient assisted living [19]. The notion of the internet-of-things comprises information acquisition and processing of all the sources outlined above, which can be considered as “sensors” in the wider sense. IoT deals with information collection and processing from virtually any type of component that can deliver observations about the surrounding environment, including both physical sensors (*i.e.*, physical devices) and virtual sensors (*e.g.*, components that process information stemming from humans, databases and/or physical devices). Hence, IoT can -on the basis of this broader definition of sensors- support all the different types of intelligence outlined above. Note also, that IoT systems are key ingredients of emerging smart cities, which include pervasive applications for smart security. This reinforces their suitability towards supporting LEAs in the task of identifying and confronting security incidents in the urban environment.

Up to date multi-sensor and IoT systems, have been extensively used in order to collect and visualize information about the surrounding environment, based on sensor information fusion and COP (Common Operational Picture) generation tools. However, (despite their suitability) they have not been used for predicting terroristic attacks. In this paper we introduce a first-of-a-kind IoT system for the prediction of terrorist attacks in urban environment. We emphasize on the presentation of the architecture of the IoT system, including a sensor information collection layer, a database for storing terroristic-related

events (so called terroristic indicators), as well as a reasoning layer aiming at processing numerous terroristic related events and identifying potential threats. For each of these layers we present an accompanying implementation. The functionalities of the overall architecture are also illustrated in the scope of a practical scenario. Furthermore, we present a set of tools for managing events over the databases associated with terroristic activities, including a web-based tool and a mobile application. The latter application is intended to be used by patrolling officers and policemen.

The rest of the paper is structured as follows: Section 2 following this introductory section presents the IoT-based architecture of the system. Section 3 is devoted to a detailed presentation of the sensor middleware that supports information collection, while Section 4 illustrates the structure of the database that underpins the system. Section 5 presents the reasoning capabilities of the system, while section 6 is the last and concluding part of the paper.

2. IoT-based System Architecture for Predicting Terroristic Attacks in Urban Environment

Our IoT-based approach to predicting terroristic attacks in urban environments focuses on proving a generic blueprint architecture for integrating terrorist prediction solutions, rather than providing an application -specific system tailored to the needs of specific scenarios. Hence, our IoT architecture has been designed as a general (reference) architecture for implementing terrorist prediction solutions, which could be flexibly customized in order to address different requirements for anticipating terroristic attacks. Thus, the architecture comprises a set of general-purpose modules that can be implemented in various ways, including different algorithms and designs. Note however that towards validating the architecture, we have instantiated the blueprint solution towards a concrete architecture implementation that comprises specific implementation modules.

A high level overview of the IoT architecture is depicted in Figure 1 **Error! Reference source not found.** At the heart of the system lies a sensor middleware framework facilitating the interfacing, collection and filtering of information from multiple underlying sensors, as well as the representation of their data into a common (standardized format). In particular, the sensor middleware framework collects information from the various sensors (including physical and virtual sensors), independently of the data format and protocol supported by the underlying sensor. Accordingly, the sensor middleware undertakes the responsibility to transform the acquired information according to a common format, based on standardized semantics. This common format ensures a unified interoperable representation of the information that stems from the multiple heterogeneous sensors. The representation is facilitated by a Geo Database, which enables the resolution of geographical information.

The IoT system targets the identification and management of terroristic indicators, *i.e.*, events that signal the possibility of a terroristic attack. These events are used in order to predict potential attacks. Therefore, an important module of the system (*i.e.*, the “event identification and formulation module”) deals with the formulation of such events. The formulation relies on the transformation of information from sensors (represented according to the standard format) to events of the Terroristic events database (which is also shown in Figure 1). The latter is designed to accommodate and persist terroristic events, including events transformed from physical and virtual sensors, as well as events provided by human operators. The latter events are entered by

human operators into the database through a GUI (Graphical User Interface) based on an appropriate mobile application.

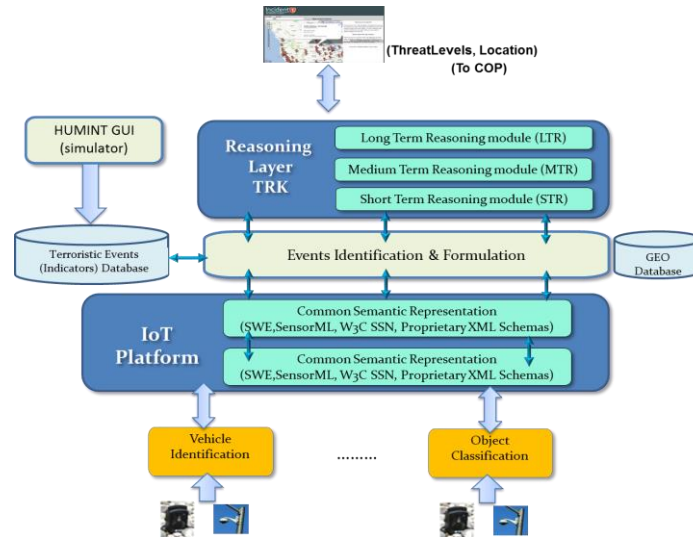


Figure 1. Building Block for IoT Architecture Boosting the Prediction of Terrorist Attacks in Urban Environments

The IoT architecture prescribes also reasoning modules, which provide the intelligence needed to infer information about the identified and anticipated threat levels associated with potential terrorist attacks. The reasoning modules are grouped in a Terrorist Reasoning Kernel (TRK) which consists of several modules operating on various levels of temporal and special scope. The TRK aims at producing reliable forecasts regarding threat levels of potential terrorist actions with special focus on requirements from different actors and users of the system. It processes events in the form of sensor-independent information (including “human” sensors) coming from the “Event Identification & Formulation” module. These events refer to suspicious/significant situations which are potential candidates for terrorist threat events within a real-time reasoning scenario. In order to efficiently process the incoming event stream, the TRK boosts its performance by exploiting histories of threats and of the related events coming from external sources over a longer time horizon. Another important part of the IoT terroristic attack prediction system is the visualization of the relevant information, such as alert levels and attack prediction levels across different timescales and within specific locations.

As already outlined we have instantiated this high-level reference architecture on the basis of specific sensor middleware implementation, reasoning modules implementation and sensors implementations. Following paragraphs illustrate the components and platforms that are implemented for the instantiation of the architecture in the scope of its validation.

3. Low-Level Information Collection and Processing

3.1. Sensor Data Collection

Our implementation of the IoT architecture is based on the open source GSN (Global Sensor Networks) middleware. This middleware provides the means for accessing contextual

data from the various sensors, including both physical devices and virtual sensors (*e.g.*, software components providing observations). GSN provides a flexible middleware for the gathering and processing of data streams generated from different sensors, based on its virtual sensor concept (VS). The latter concept enables different sensors to be described and integrated in the middleware based on a common XML format. Following the description of sensors as virtual sensors and their deployment in the GSN middleware, the GSN platform supports: (i) data acquisition from various sensors, (ii) filtering of data based on an intuitive, enriched SQL syntax, (iii) execution of customizable algorithms on the results of the query, and (iv) output of the generated data based on its notification subsystem. Overall, the exploitation of the GSN sensor middleware ensures:

- Support for virtually any type of sensor and data stream, through minimal effort. Sensor networks and data streams can be specified in a declarative way using XML as the syntactic framework and SQL as the data manipulation language
- Flexibly addition of new types of sensor networks, along with dynamic (re-) configuration of the system during run-time without having to interrupt on-going system operation.
- Support for very large numbers of data producers and consumers with a variety of application requirements.
- The provision of easy-to-use web-based management tools for the deployment and configuration of the sensors.

The sensor middleware provides the means for different sensor components (*e.g.*, contextual processing algorithms, signal processing algorithms) to provide/transmit their data to the IoT system according to a common set of metadata, which is specified in a common data feed specification. The latter specification is implemented as an XML/JSON schema. The specification defines a general entity for describing feeds. Each feed consists of the feed description and a set of Components and Outputs. It also includes a title and a textual description of the function it performs. Furthermore, every feed is characterized as “Physical” or “Virtual”, depending on the components that comprise the feed. Physical sensors are associated with physical sensing devices, while virtual sensors correspond to software elements that produce observations. Among the obligatory properties of a feed are the descriptions of the components attached to the feed, as well as the outputs it produces. Furthermore, each feed has a small set of optional, but commonly used, properties which include: geo-location information, descriptive tags and contact information. Also, the feed specification allows the description of a feed’s outputs using a type for describing their properties, such as the data type of the retrieved measurements and the Measurement Unit (Optional). Overall, the data feed specification (as described in the corresponding XML schema) allows providers of sensing components to formulate their feeds in a way that could make them usable in the scope of the terroristic prediction system.

3.2. Sensor Data Processing

The sensor middleware layer of the system enables also the processing, combination and fusing of multiple data feeds, stemming from heterogeneous sources. In particular, the middleware enables the definition of new virtual data feeds that combine information from two or more data feeds. In this way it allows for the implementation of multi-source data fusion, giving rise to the support of the JDL model for data fusion. Note that the fusion of multiple data sources can be supported via the definition of new virtual sensors based on the SQL-like language of GSN, but also through the integration of data processing frameworks (*e.g.*, rule engines, machine learning frameworks) over the various data feeds. The ultimate goal of the data processing at the sensor middleware level is to ensure the formulation and

integration of terroristic-attack related events into the terroristic events database, which is illustrated in the following section.

4. Terroristic Events Database and Data Management Interfaces

4.1. Database Modeling and Terroristic Indicators

While the use of the sensor middleware ensures the IoT system's capabilities towards accessing, collecting, transforming and fusing heterogeneous sensor streams, the terroristic events database ensures that events identified through processing multiple sensors and observations are later persisted according to a structure that enables their use for reasoning about potential terroristic attacks and actions. Since the IoT system has to be generic and usable for detecting terroristic attacks across multiple scenarios, the database design should take into account the most common known activities that when observed they can signal indications of terroristic attacks. To this end, the database comprises events and semantics that if promptly detected and analyzed, could provide information to prevent future terrorist attacks. Indeed, the design of the PROACTIVE database was based on the identification and documentation of semantics (*i.e.*, events and activities) that are known to be related to the planning and organization of terroristic attacks.

Towards identifying and structuring the (general scenario-agnostic) semantics of terroristic attacks, we studied relevant works, which describe the most common indicators of terroristic attacks [3][13]. These works provide insights on the data and metadata that should be identified and processed by the IoT system. In several cases they converge on a common set of data and metadata associated with the events that could indicate preparation or anticipation of terroristic attacks. Therefore they provide a sound basis for modeling data in the scope of the PROACTIVE database. According to [3], there are eight families of indicators of future terroristic attacks and relevant security incidents. These families indicate a potential taxonomy of terroristic events, which can be used to classify the various contextual cues that will be identified by the IoT system. The following table illustrates the eight different categories of terroristic indicators, along with some sample set of indicators for each category.

Table 1. Classification of Terrorist Indicators and Examples

Terrorist Indicator Category	Sample Indicators (examples)
Preoperational Surveillance	<ul style="list-style-type: none"> • Foot surveillance involving two or three individuals working together. • Mobile surveillance using bicycles, scooters, motorcycles, sport utility vehicles, cars, trucks, boats or small aircraft. • Persons or vehicles being seen in the same location on multiple occasions. • People sitting in parked cars for an extended period of time.
Seeking and Eliciting Information	<ul style="list-style-type: none"> • Inquiries about size of security force. • Inquiries concerning access to sensitive areas.

	<ul style="list-style-type: none"> • Inquiries regarding the licensing/certification for hazardous materials transportation.
Probing and Testing Security Measures	<ul style="list-style-type: none"> • Initiation of false alarms (e.g., a bomb threat). • Attempts to penetrate physical security barriers. • Attempts to test physical security/response procedures at key facilities. • Attempting to get weapons or other restricted materials through a security checkpoint, such as a metal detector or bag search point.
Intrusion (against physical security or cybersecurity measures)	<ul style="list-style-type: none"> • An intruder enters a restricted area with malicious intent, damaging or manipulating some system of the target. • Intrusion into a computer network. • Unauthorized personnel entering a restricted area for the purpose of collecting information or stealing something associated with a target.
Acquiring Supplies	<ul style="list-style-type: none"> • Suspicious or improper attempts to acquire official vehicles, uniforms, badges, access cards, or identification for key facilities • Theft of two-way radios or scanners. • Theft or purchase of paint or logos similar to those found on security or emergency vehicles.
Identification of Suspicious People	<ul style="list-style-type: none"> • Persons or vehicles observed in the same location on multiple occasions and/or those who engage in unusual behavior • Persons observed near a potential target using or carrying video, still camera, or visual enhancement devices (telescopes, binoculars, night vision goggles). • Persons showing an interest in or photographing the security measures at a target • Persons drawing pictures or taking notes in a non-tourist area not normally known to have such activity.
Dry Run or Trial Run of an Attack	<ul style="list-style-type: none"> • Suspicious persons sitting in a parked car for an extended period of time for no apparent reason. • Persons observed monitoring a police radiofrequency and recording emergency response times • Photography or videotaping with no apparent reason. • Abandoning object(s), such as pieces of luggage.
Deploying Assets and Getting in Position	<ul style="list-style-type: none"> • Loading Weapons and other supplies in vehicles • Suspicious Behaviors • Deployment of weapons

Each of the indicators listed in the above table can be captured through either sensing devices or human sensors. In order to record and keep track of these events within the IoT system, the terroristic database has been designed in a way that associates the various indicators with their category, but also with the sensor (*e.g.*, device, system and algorithm, human) that reported the event. Furthermore, the database keeps additional information required by the reasoning layer, such as the date and time of incident, the specific location of an incident, a description of the incident, a description of the facility or person being targeted, the number of adversaries that were conducting the surveillance and more. Data management within the terroristic events database is empowered by two applications which are described in the following paragraph.

4.2. Data Management Applications

Two different data management applications are provided in order to manage data within the terroristic events database. These applications are tools that accompany the instantiation of the IoT architecture and include: (a) a data management application featuring a ubiquitous web based interface and (b) a mobile phone interface (mobile app) implemented for the mainstream Android devices. Primarily, these applications enable users of LEA's located in urban environments, in external or internal locations, to report location-based events. The web based application interface is primarily modeled for administrative use (*e.g.*, as part of the on-premise infrastructures) while the mobile application is intended for use by patrolling officers.

Key functionalities of both implemented applications include:

- The ability to select and report different event types, in different locations. The applications automatically select the terroristic events/indicators that pertain to specific locations.
- Location-based reporting leveraging the GPS capabilities of the users' devices in order to accurately report the position of the event.
- Instant Event Reporting, through a minimal number of interactions (clicks) to the interface. The interactions concern the selection of the event type and its location.
- Customized Event Reporting, which allow the user to indicate/report new customized events on the basis of a custom location field and a customized event description.

Figure 2 **Error! Reference source not found.** Illustrates the user interface of the web application. It provides a map mashup, which allows users to locate events on the map prior to reporting them. The web application provides the means for selecting event types corresponding to terroristic indicators in the terroristic events database. The interface provides the locations of all virtual sensors on a map in order to facilitate the use of the applications at the LEA's headquarters/premises.

The screenshot shows a web-based application interface with two main panels. The left panel, titled "Sensor data...", contains a form for entering location information. It includes a "Location Category" dropdown menu set to "scenario_a_locations", a "Location" dropdown menu set to "ministry_of_finance", and "Coordinates" input fields for "Lat" (37.976538) and "Lon" (23.733277). Below these fields is a "Submit" button. The right panel, titled "Sensor list", displays a table with two columns: "ID" and "Description". The table contains two rows: ID 6 with description "civilian" and ID 7 with description "police car". Below the table are "Event" and "View" buttons. A map of Greece is visible in the background of the left panel.

ID	Description
6	civilian
7	police car

Figure 2. Snapshot of the Implemented Web based Application

Similarly, the mobile application provides “on-the-go” capabilities for use by police/security patrols or even by civilians. It also provides a uniform way of reporting events, based on the indicators and event types of the database (as shown **Error! Reference source not found.**).The application targets the Level 17 Android API (Jelly Bean), while it is compatible with Level 10 Android API and above (Gingerbread). In this way it is available on over 97% of mobile devices running the Android platforms.

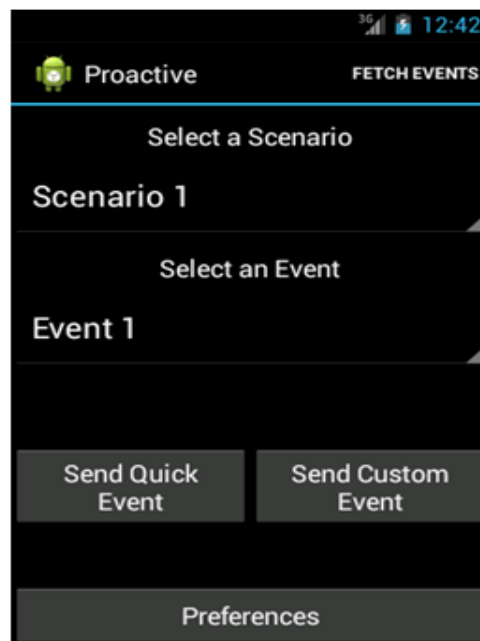


Figure 3. Snapshot of the Mobile App Data Management Application Interface

5. Reasoning Layer

There has been an increased research and development activity in the terrorist attack/threat detection/prevention domain. A number of domain specific computational approaches have

already been proposed for improving both data collection and data analysis. The overall aim is to forecast long term activities of terrorist groups or to suggest prevention by analyzing terrorist group past behaviors.

On one side we approach the design of the reasoning layer by taking into account both the abstraction levels of the potential information sources (sensor information, police patrol inputs, news events, external semantic crafted data sources) and the expert user roles that are currently defined as crucial in the intelligence analyst flow for analyzing/detecting potential terrorist threats. Hence three modules can be identified as: the Short Term Reasoning (STR) module, processing symbolic events generated by virtual sensors and considering only short-term time horizon, the Medium Term Reasoning (MTR) module operating on a larger time horizon, and the Long Term Reasoning (LTR) module that takes into account histories of threat events and/or external data sources spanning over wider time and spatial horizons.

Taking into account the aspect of various expert user roles of our design scenario we note that the STR module works at the level of event generator users, that is police patrols, that represent “human” virtual sensors and can notify the system in case of the occurrence of an suspicious situation taking into account short histories of events. The MTR module works at a level of tactical user that intervene in order to indicate the sensitivity/criticality of the current situation and the LTR module works at the level of strategic users taking into account a longer scenario. All of the above mentioned modules consider user feedback by interpreting end user actions that can be obtained through the COP interface (Figure 1).

On the other hand we approached the design of the reasoning layer with having in mind the functional requirements of the overall IoT terrorist attack prediction system. As the system aims at producing timely proactive responses by relying on reliable forecasts about terrorists’ actions, there is a direct correspondence to advanced Complex Event Processing (CEP) systems. Examples of CEP applications include environmental monitoring, monitoring for transportation and logistics, trading for financial markets, security application for intrusion detection, bio-hazard attacks etc. Events are either defined as primitive (atomic) events (*e.g.*, sensor data, trading ticks, credit card transaction, network signals) and complex (composite) events (*e.g.*, landslide, terrorist attack, plane landing, and credit card fraud). Overall CEP systems (or as defined in [5] information flow processing systems) are able to manage multiple data stream sources and derive new information about the data stream through use of a set of processing rules.

The IoT terrorist attack prediction system has two major requirements, the first one being real-time delivery of relevant information in a useful format (*e.g.*, outcomes of end user workshops with domain experts indicated a timely response with maximum latency of 1 minute), a requirement that can be adequately approached through a CEP framework. And the second major requirement being proactively preventing events before they occur and not only reacting after they happen, specifically relevant in applications as credit card fraud, terrorist threat prevention *etc.*, The value of the detected complex event decreases with time as described in [7] and clearly there is a higher importance of reacting in near real-time to terrorist threats notifications as opposed to a day after and even more so in proactively detecting/predicting terrorist threats ahead of time. To address this issue we jointly consider both CEP systems and predictive analytics approaches, in this way enabling processing online streams of events while and inferring decision/ future events of inters based on past and current events. In [20] a discriminative analysis is used to detect a suspicious combination of events from an event cloud of an organization and in [7] classical CEP system is used to generate training sets and model predictions with decision trees.

We adapt the approach of using statistical machine learning for discovering patterns in incoming event streams, since the event types coming into the IoT terrorist attack prediction

system can be both from multiple sensors (including a “human” sensor) and numerous types (detected actions) thus, crafting rules in a terrorist scenario, can become a burdensome process even for an expert. Statistical modeling of sets of suspicious actions into a threat event enables “soft” rules that adapt with time and are able to follow the shift in behavior of terrorist groups. Furthermore, we maintain our designed decisions regarding above mentioned abstraction levels of the information sources and mapping to the expert user roles. Hence, predictions rules are learned both long-term and short-term historical data and integrated with rule based CEP systems through STR, LTR and MTR module of the TRK component (the reasoning layer). The above mentioned combination of both declarative and soft rules discovered through machine learning is particularly useful in applications where a certain level of uncertainty regarding complex events is allowed, as in the IoT terrorist attack prediction system.

In this section we describe our implementation of the reasoning layer through the TRK component, specifically the separation into short, medium and long term reasoning modules. This division was based on the various types of data sources these modules can process, and on the temporal and special scope of the data it can reason over. The individual data modules and data flows can be observed in Figure 4 and are elaborated in the follow-up.

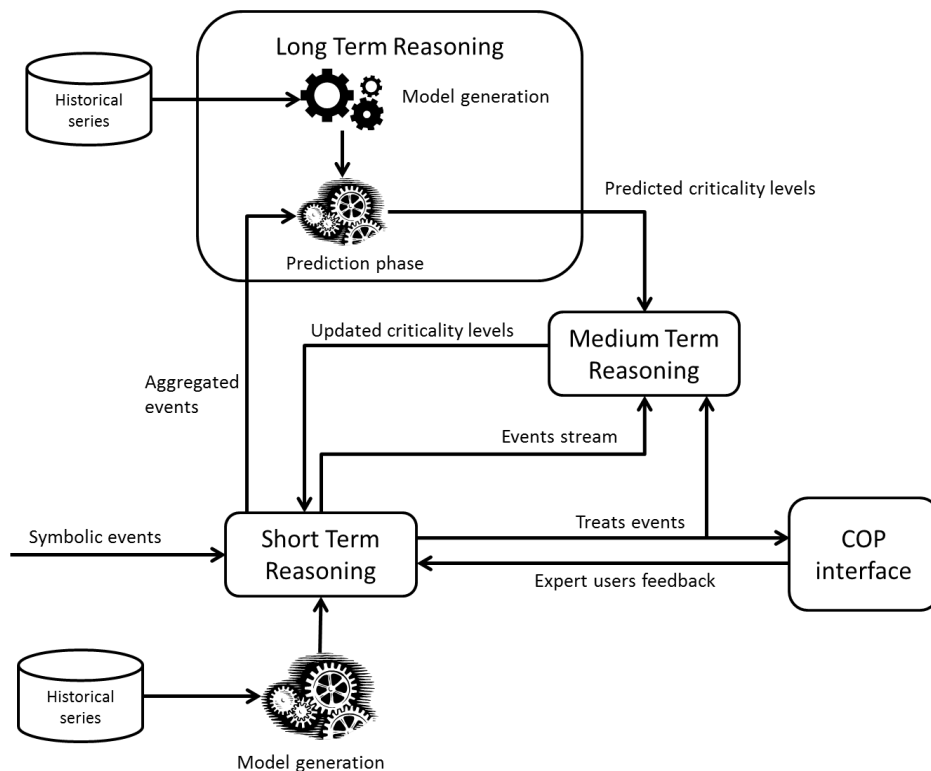


Figure 4. The TRK Modules and Data Flows

5.1. Short-Term Reasoning

Since intelligence analyst users are typically required to try to identify an attack from within thousands of alerts coming from numerous surveillance systems, Situation Awareness (SAW) systems are of high importance in large control centers (e.g., air and road traffic management). Their goal is to reduce the information overload of operators induced by

various data sources, as they risk lack of situation awareness due to limited abilities of existing systems focused on mere presentation of available information [2]. Hence we focus on a capability to aggregate alerts together but also to capture patterns that can detect suspicious situations, therefore reducing the amount and improving the relevance of information that an analyst needs to consider. As the role of the STR module is real time threats detection, based on events within a short time-window, we rely on classical Hidden Markov Models, able to model sequences of events that include and represent a kind of partial history of recent occurrences, and capturing the relevant ones that might indicate a threat [4]. The STR module considers symbolic events in the form of sensor-independent information (including “human” sensors) about suspicious /significant situations within a real-time reasoning scenario. It is built upon the idea of micro-environments, which are associated with a physical environment of limited size and complexity (*e.g.*, a building with its surroundings). Additionally, the STR module has a built-in relational taxonomy regarding relevance of symbolic sensor-independent events from the incoming event stream (output of the “Event Identification & Formulation” module) in a specific context of a micro-environment (*e.g.*, a car being parked in a parking lot of a commercial center may be of low relevance, whereas the same event in front of a Ministry building may be of medium or high relevance). The relevance (*e.g.*, Not Relevant, Low, Medium, and High) is defined by an expert user upon system deployment and depends on the type of local microenvironment being considered (*e.g.*, building, square, metro station).

A micro-environment captures the required sensitivity of the modeled environment and hence corresponding criticality levels through a threshold parameter that is used for detecting a threat event. The need for different sensitivity settings is an outcome of an end user workshop with domain experts that noted different levels of alertness to events and changes to their own decision making process in correlation to the pre-estimated criticality of the situation at hand. The end goal of the micro-environment is to compute the probability of a threat in the incoming event stream in the context of the previous state, that is, the criticality level of the micro-environment and the relevance of the observed suspicious event.

During run-time the current criticality levels are defined either by a COP user (domain expert users) or set/modified by medium and long-term reasoning components. In the initial setting three levels of criticality (Low, Medium, and High) are modeled with a HMM where the selection of appropriate threshold parameter depends on the current run-time criticality of a particular micro-environment. The criticality levels of a microenvironment define different sensitivity to incoming suspicious events of a particular physical environment (*e.g.* in a high criticality setting an occurrence of a suspicious event raises the probability of threat faster than in a lower criticality setting) see Table 3 and Figure 5.

Furthermore the STR module includes the definition and training of different micro-environments based on histories of events, where the initial parameters are estimated from simulated data from a game playing scenario by experienced domain users (police versus terrorists) and from the initial storyline defined through threat event indicators stored in the database.

Figure 5 depicts different behavior of the STR for two criticality level settings (low and high) for one low relevance incoming event (on the left hand side) and many low relevance incoming events (on the right hand side). The left hand figure shows effects of a single low relevance event in a high criticality level setting that can easily overpass the threshold triggering a threat alert (here for visualization purposes kept at value 0.5), while in a low criticality level setting the probability of threat will only slightly increase. The right-hand side of Figure 5 demonstrates a case of multiple low relevance events that in the high criticality

level setting maintain the high probability of a threat for a longer time while in a low criticality level setting influence a gradual increase of the probability a threat. Similar behavior can be learned from training data capture for events with different levels of relevance.




Figure 5. The STR Behavior for two Criticality Level Settings (Low and High) for One Incoming Low Relevance Event (Left Hand Side) or Many Incoming Low Relevance Events (Right Hand Side)

5.2. Long-Term Reasoning

Following the ideas behind [11], [21] we design an approach that exploits correlations between action and surrounding context features of a representation of terrorist attacks and threats. The mentioned feature design fits our system as events detected by the “Event Identification & Formulation” module from Figure 1 representing potential terrorist attraction that could lead to a threat and the surrounding information as context, particular location and time *etc.*, This design feature was specifically considered in the long term reasoning (LTR) module which aims at exploiting information from both external databases, containing histories of events that represent terrorist attacks and histories of events/threats from the IoT terrorist attack prediction system itself.

The Long-term reasoning (LTR) module generates predictions of criticality levels for specific types of micro-environments from the STR that correspond to types of physical environments (*e.g.*, public buildings, metro stations, and so on). These activities are performed by analyzing histories of threats and their associated events on a long time horizon. Its ability to “predict”/re-estimate the context of the systems relies on external sources such as intelligence scenarios and/or scheduled activities (*e.g.*, planned high-profile events like “visit of a minister”). This approach allows us to group instances (either histories of threat events or terrorist attacks from external sources) from the selected dataset into clusters that ideally capture models of similar threat events (from usage history data) and/or of attacks (from external sources, *e.g.*, GTD [10], MAROB [3] datasets). The instances are represented as multi-dimensional feature vectors comprised of action features which are mapped to incoming events and of context features which are mapped to types of physical environments.

The action feature space is used as bases for building statistical cluster models and capturing the underlying similarity. Furthermore, as for each instance, the set of action features is linked to a corresponding set of context features, we capture the distribution of different physical environment context features across the clustered action feature space. In this way we model the correlation between combinations of action events and corresponding physical environments. Through building these models we aim at capturing the “typical behavior” of a terrorist attack that can potentially lead to a threat and the related physical environment. In the operational run-time stage new events are introduced in the system generating histories of threat/no threat events to be used. The

reasoning phase of this, takes as an input aggregated event capturing the short-term histories of events across a time window of analyses and position it in the closest cluster defined in the action feature space. Hence, by exploiting the correlation between action and context space the presence of each type of physical environments in the identified closest cluster can be obtained. This leads to the probability of a threat event happening in each type of physical environment that represents the sensitivity (criticality value) for each type of micro-environment in the STR module. As depicted in Figure 4, the LTR module can be divided into two different phases: the model generator phase that learns the cluster models and the long-term prediction phase that uses these models to infer the criticality levels of each type of physical environments.

5.3. Medium-Term Reasoning

The Medium-Term Reasoning (MTR) module supports overall decision making by re-evaluating the criticality level within a larger temporal scope than STR module and shorter temporal scope than LTR module. It considers the used criticality level (STR), the predicted criticality level (LTR), the short-term histories of symbolic and threat events, and feedback, when that is available, from domain expert users regarding the criticality level (COP interface) see Figure 4. Since the role of MTR reasoning module in the TRK is to process numerous incoming events types of different nature and update the sensitivity of the STR in near real-time manner, we believe that event processing approaches fits requirements of the MTR module. Specifically we consider CEP systems [5] as they filter and combine incoming events of particular patterns from the external world to understand what high-level complex events have occurred and notify relevant actors or be reused as an input in the CEP solution. For the implementation of the MTR module we used Drools¹, an open source project. More specifically, the components used are the rule engine and the module that enables event processing capabilities. This framework is used to build expert systems, as it is a declarative rule based coding environment for defining, executing and maintaining rules. Domain knowledge is to define a set of declarative rules of type:

***when:** < UCL:= Low and PCL:=Medium and over window : time (N) threat events are more than Threshold % of all events>*

***then :** <CL:= Medium>*

This type of rules are define for all combinations of input parameters (Used Criticality Level (UCL), Predicted Criticality Level (PCL)) and output value (updated Criticality Level(CL)) as well as take into account if the UCL is set by an expert user or by the previous phase of MTR module.

5.4. TRK Application

In order to demonstrate the functionalities of TRK modules, as for the data management application, a prototype has been developed through which it is possible to monitor threats probability and information of each micro-environment through specific interfaces. In this way, the prototype interface displayed in Figure 6 allows strategic users to monitor threat probabilities for different micro-environments (e.g., square, building, and metro station) and change the criticality level for different instances of micro-environments. The type of the micro-environment for every initiated instance can be set during configuration phase. The COP interface is used for presentation of the detected threats to the user and displays how many threats have been detected so far and how many of them need user attention.

¹<http://docs.jboss.org/drools/release/6.0.1.Final/drools-docs/html/index.html>

Additionally, for each event, the information regarding a specific threat, as time stamp, location, source, duration, object class, action, features and criticality level is available.

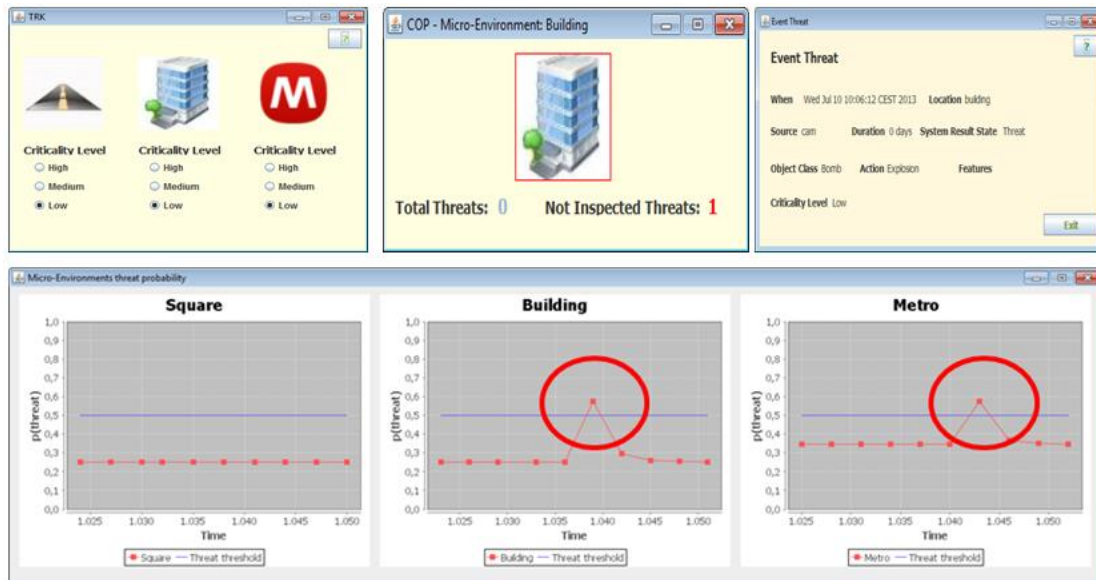


Figure 6. TRK Interface for Setting Criticality Levels of a Micro-Environment (Top-Left Figure), COP Threat Presentation Interface (Top-Middle Figure) and COP Event Inspection Interface (Top-Right Figure). The Bottom Figure Displays the Threat Probability Monitoring Interface for the Case of a Threat Detected for in Two out of Three Different Types of Micro-Environment

For demonstration and simulation purposes we also provide an Events Generator (EG) interfaces for creation of events that are processed by our micro-environments. Each micro-environment can have multiple numbers of EG interfaces depending of the number of virtual sensor that are related to it (*e.g.*, Police Patrol and Camera as it is displayed in Figure 7).

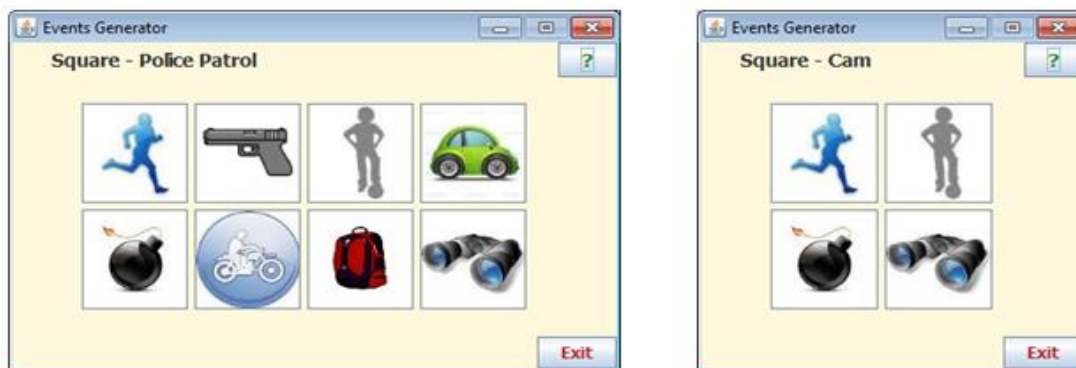


Figure 7. Events Generator (EG) Interface for “Square” Type of Micro-Environment and Two Virtual Sensors: Police Patrol (Left-Hand Side) and Camera (Right-Hand Side)

6. Validating Scenario

6.1. Overview

Towards the validation of our IoT architecture, we have devised and implemented several terrorist prediction scenarios. In the sequel we presented an integrated validation scenario that is supported by the implemented system. This scenario distinguishes two different phase/stages where events are collected:

- A pre-operational stage, where events (associated with terroristic indicators) are entered in the database. This stage involves a coarse timescale of weeks or even months associated with the collection of these events.
- A “real-time” operational phase, where the IoT system identifies specific live events and signals terrorist attack prediction alerts. To this end, the system exploits its reasoning capabilities over events collected/persisted within the database.

In the remaining section we demonstrate the work flow of the IoT system, and the ability to capture possible terrorist attacks.

6.2. Pre-Operational (Off-Line) Phase

This phase extends itself in a period of four weeks (28 days) before the commencement of the real-time phase. The following table lists the pre-operational events/indicators observed and collected throughout this time interval using the IoT system:

Table 2. Events Collected During the Pre-Operational Stage of the Demonstration Scenario

Pre-operational Event / Terrorist Indicator	Location(s)	Time(s)	Sensor Capturing the Event
Vehicle (with license plate YYY-XXXX) Identified outside the Building of the Ministry of Finance Multiple Times	Building of the Ministry of Finance	Day #1, 20:00 Day #6, 15:00 Day #11, 23:00	Vehicle Tracking, License Plate Identification Component (WP7)
Persons Sitting in Parked Cars for more than 1hour Multiple Times	Building of the Ministry of Finance	Day #3, 17:15 Day #12, 21:23	Human Sensors / Police Patrols
Persons Sitting in Parked Cars for more than 1hour Multiple Times	Parliament Building	Day #5, 20:32 Day #14, 22:45	Human Sensors / Police Patrols
Suspicious Person Photographing the area	Building of the Ministry of Finance	Day #3, 10:15 Day #6, 14:12 Day #9, 12:15	Human Sensors / Police Patrols
Suspicious Person Photographing the area	Conference Center	Day #6, 15:32 Day #8, 13:45	Human Sensors / Police Patrols
Inquiries concerning access to sensitive areas	Building of the Ministry of Finance	Day #12, 11:45	Human Sensor / Asked via Telephone
False (Security) Alarm	Conference Center	Day #14, 13:12	Human Sensor
Network Intrusion attempt	Ministry of Finance Network	Day #14, 17:45 Day #16, 3:34	Human Sensor / Reported by the ICT Security Officer

A Gun was Stolen	N/A	Day #1, 10:35	Human Sensor / Reported to the Police – Entered to the PROACTIVE database
Person Loitering	Conference Center	Day #18, 10:32 Day #21, 10:28	PROACTIVE Person Tracking Technology (WP7) and Human Sensor (Patrol)
Person Loitering	Building of the Ministry of Finance	Day #17, 13:10 Day #17, 13:10 Day #26, 9:41	PROACTIVE Person Tracking Technology (WP7) and Human Sensor (Patrol)
Person taking notes	Parliament Building	Day #15, 14:12 Day #17, 15:15	Human Sensor (Patrolling Policeman)
Person taking notes	Building of the Ministry of Finance	Day #12, 11:41 Day #20, 18:15 Day #27, 16:12	Human Sensor (Patrolling Policeman)
Person taking notes	Conference Center	Day #12, 9:21	Human Sensor (Patrolling Policeman)
Person observed with facility notes and a map	Building of the Ministry of Finance	Day #18, 10:31 Day #28, 14:42	Human Sensor (Patrolling Policeman)
Person observed with facility notes and a map	Conference Center	Day #13, 12:32	Human Sensor (Patrolling Policeman)

These events are observed either through sensing systems or by human sensors such as patrolling policemen. For the purpose of the demonstration, these events were entered and persisted in the database through the data management applications illustrated in the previous sections. However, during a production use and operation of the system, these events can be entered into the system either by the interfaces developed for human sensors (such as the Mobile app for patrolling policemen), or through the sensor middleware infrastructure of the architecture.

6.3. Real-Time Phase

The real-time phase of the project comprises events that will be captured during the real-time operation of the IoT system. For demonstration purposes, as part of this phase real-life sensors and sensor processing systems are deployed, along with data management interfaces for human actors. For purpose of the demonstration the EG interface from Figure 7 is being used to simulate the input stream of suspicious events coming from either a police patrol or a camera. The following table denotes several types of suspicious events captured during the real-time operation of the system and whether they trigger threat notifications or not after being processed by the STR module in function of the criticality levels detected/modified by LTR and MTR modules.

It is assumed that the real-time operation of the system takes place during Day#29 and takes into account knowledge gathered by analyzing sequences of symbolic events (produced by the virtual sensors of the IoT system) that correspond to one of the events collected during the pre-operational stage of the demonstration scenario (Table 2). The pre-operational events correspond to examples of terrorist indicators and were generated respecting the eight categories of indicators of future terroristic attacks, as shown in Table 1 **Error! Reference source not found..** The demonstration scenario outlines pre-operational events that have taken place during the previous four weeks at

the assets/locations. In order to successfully integrate the long-term demonstration scenario, on one side the TRK models sequences of symbolic events associated with the pre-operational terrorist indicators through its STR module, and on the other side gathers associated terrorist indicators categories through its LTR module. The overall goal is to exploit longer –term scenarios and potentially fill gaps in the puzzle concerning a possible terrorist attack. Therefore two relevant mappings devised by domain experts are used, a mapping between examples of terrorist indicator events in a demo scenario (first column of Table 2) and the terrorist indicator categories (first column of Table 1), and a mapping between terrorist indicator events in a demo scenario (first column of Table 2) and the symbolic sensor events (first column of Table 3).

Table 3. Example of Events Identified During the Real-Time Stage of the Proof-of-Concept Demonstration Scenario

Short term Event / Terrorist Indicator	Location(s) (micro-environment)	Criticality Level of the micro-environment	Sensor Capturing the Event	Trigger in the PROACTIVE System
Unusual Movement of Vehicle	Building of the Ministry of Finance	Low	PROACTIVE Unusual Motion Detector Technology	No threat notification.
Unusual Movement of Vehicle	Building of the Ministry of Finance	Medium	PROACTIVE Unusual Motion Detector Technology	No threat notification for a single event. Threat notification for frequently repeated events.
Unusual Movement of Vehicle	Building of the Ministry of Finance	High	PROACTIVE Unusual Motion Detector Technology	Threat notification even for single event.
Unusual Movement of Vehicle	Parliament Building	Low	PROACTIVE Unusual Motion Detector Technology	No threat notification. Threat notification for frequently repeated events.
Unusual Movement of Vehicle	Parliament Building	Medium	PROACTIVE Unusual Motion Detector Technology	Threat notification even for single event.
Unusual Movement of Vehicle	Parliament Building	High	PROACTIVE Unusual Motion Detector Technology	Threat notification even for single event.
Observation of Parked Car for an extended period of time for no apparent person	Building of the Ministry of Finance	Low	Human Sensor – Police Patrol (Mobile App)	Threat notification even for single event.

Observation of Parked Car for an extended period of time for no apparent person	Building of the Ministry of Finance	Medium	Human Sensor – Police Patrol (Mobile App)	Threat notification even for single event.
Observation of Parked Car for an extended period of time for no apparent person	Building of the Ministry of Finance	High	Human Sensor – Police Patrol (Mobile App)	Threat notification even for single event.

In the Table outlined above, the changes to the threat levels are produced by the TRK component, which decides the particular threat levels and alarm levels to be displayed in the user interface (*i.e.*, COP) of the system Figure 7.

Additional to the EG interface a practical simulation of the real-time phase is implemented in a game playing scenario, where several actors are involved including patrolling policemen (*e.g.*, 2-3 policemen), security forces officers (*e.g.*, 4-5 officers) deployed at specific assets to be protected, drivers or owners of parked cars, as well as actors playing the role of civilians in near the assets to be protected by the system. This information can be used to train components of the TRK, specifically to build the initial models for the STR module. In such a setting, the IoT system can also serve as a training tool for officers, policemen and other LEA's employees help them more easily access and respond to potential terrorist attacks.

7. Conclusions

In this paper we have outlined the main architectural and structural principles that underpin the implementation of an IoT-based system for the prediction and anticipation of terrorist attacks. The introduced architecture can serve as a blueprint for the implementation of sensor-based systems for the prevention of terroristic attacks in urban environments. Towards a practical implementation and validation of the introduced architecture, we have implemented a wide range of software and middleware components, including:

- A sensor middleware framework, which provides the means for collecting and processing data from multiple sensors, while at the same time supporting the implementation of a wide range of information fusion and reasoning algorithms, which form the basis for predicting terroristic attacks. The sensor middleware facilitates the interfacing to multiple heterogeneous sensors and information sources, while also supporting the representation of the various data streams and data structures in standards-based formats.
- The reasoning layer (TRK component), composed by three sub modules (STR, MTR and LTR), which produces forecasts regarding terrorist attacks by carefully taking into account the abstraction levels of the potential information sources, the expert user roles and the functional requirements of the overall IoT terrorist attack prediction system. In particular, the Short Term Reasoning (STR) module processes symbolic events generate by virtual sensors and considering only short-term time horizon, the Medium Term Reasoning (MTR) module operates on a larger time horizon, and the Long Term Reasoning (LTR) module that takes into account histories of threat events and/or external data courses spanning over wider time and spatial horizons. The TRK comes with a prototype that allows expert users to monitor threat probabilities for

different micro-environments, change their criticality level, display the detected threats and their specific information (*e.g.*, time stamp, location, source, *etc.*).

The introduced IoT architecture can be enhanced and flexibly customized to more specific business requirements and scenarios. In this sense, it acts as a blueprint for the implementation of terrorist prediction systems in urban environments. The implementation of such system may entail the enrichment of the core system with additional sensors, databases, middleware components and reasoning algorithms.

In order to validate the architecture (in terms of its functional characteristics) we have illustrated an integrated scenario, which involves the main components of the architecture and is based on the identification and processing of some events that are commonly associated with terroristic attacks. This integrated scenario has served as a basis for validating the prototype implementation of the various components of the architecture. The validation of the system based on additional scenarios will serve as a basis for fine-tuning the implementation of individual components, but also for benchmarking and auditing the effectiveness and accuracy of the produced alerts and notifications.

8. Related Works

Following is a brief description of related papers/studies:

- **A Simple Ontology for the Analysis of Terrorist Attacks**
This paper describes a foundation for an Ontology that represents terrorist groups and their intentions and the entities associated with such an environment. This ontology attempts to generalize the organization and classification these entities, providing a preliminary foundation for a larger system [22].
- **The prediction of terrorist threat on the basis of semantic association acquisition and complex network evolution**
This paper is a Continuation of an MCC 2006 Conf. publication, where various methods are discussed in the context of Early Warning Systems for terroristic action preparation activities. These methods rely on semantic and complex networks that are used to extract relevant information [23].
- **Indicators of Terrorist Activity – Stopping the Next Attack in the Planning Stages**
This is an analysis that identifies the various terroristic preparatory actions that may indicate potential future attacks. Such actions may include theft or purchasing of vehicles, materials and even abduction of persons that may be in any way exploited in performing terroristic attacks[13].
- **TRAKS: Terrorist Related Assessment using Knowledge Similarity**
This project proposes an early warning system that detects money laundering, terrorist planning, and id theft through knowledge similarity[24].

Acknowledgements

Part of this work has been carried out in the scope of the PROACTIVE project (FP7-285320) (<http://www.proactive-project.eu>). The authors acknowledge help and contributions from all partners of the project.

References

- [1] K. Aberer, M. Hauswirth and A. Salehi, “Infrastructure for Data Processing in Large-Scale Interconnected Sensor Networks”, MDM, (2007), pp. 198-205.

- [2] N. Baumgartner, W. Gottesheim, S. Mitsch, W. Retschitzegger and W. Schwinger, "Be Aware Situation awareness, the ontology-driven way, Data & Knowledge Engineering, vol. 69, Issue 11, (2010), pp. 1181-1193.
- [3] B. T. Bennett, "Understanding, Assessing and Responding to Terrorism", Protecting Critical Infrastructure and Personnel John Wiley & Sons, Inc., (2008).
- [4] C. C. Chen, M. C. Chen and M. Chen, "An adaptive threshold framework for event detection using HMM-based life profiles", ACM Trans. Inf. Syst., vol. 27, Article 9, (2009), pp. 2.
- [5] G. Cugola and A. Margara, "Processing flows of information: From data stream to complex event processing", ACM Comput. Surv., vol. 44, no. 3, (2012), pp. 15.
- [6] C. Doulaverakis, N. Konstantinou, T. Knape, I. Kompatsiaris and J. Soldatos, "An Approach to Intelligent Information Fusion in Sensor Saturated Urban Environments", EISIC, (2011), pp. 108-115.
- [7] L. J. Fulop, A. Beszedes, G. Toth, H. Demeter, L. Vidacs and L. Farkas, "Predictive complex event processing: a conceptual framework for combining complex event processing and predictive analytics", In Proceedings of the BCI. ACM, New York, NY, USA, (2012), pp. 26-31.
- [8] M. Hecking, "Content Analysis of HUMINT Reports", In Proc. of the 2006 Command and Control Research and Technology Symposium (CCRTS) "THE STATE OF THE ART AND THE STATE OF THE PRACTICE", San Diego, California, (2006).
- [9] G. Hummel, M. Russ, P. Stütz, J. Soldatos, L. Rossi, T. Knape, Á. Utasi, L. Kovács, T. Szirányi, C. Doulaverakis and I. Kompatsiaris, "Intelligent Multi Sensor Fusion System for Advanced Situation Awareness in Urban Environments", Future Security, (2012), pp. 93-104.
- [10] G. La Freea and L. Dugana, "Introducing the Global Terrorism Database, Terrorism and Political Violence", vol. 19, (2007), pp. 181-204.
- [11] V. Martinez, G. I. Simari, A. Sliva and V. S. Subrahmanian, "Convex: Similarity-based algorithms for forecasting group behavior", IEEE Intelligent Systems, vol. 23, no. 4, (2008), pp. 51-57.
- [12] M. J. Licata, "Citizens Terrorism Awareness and Survival Manual (C.A.T. Eyes)", (2008).
- [13] "RISS/ROCIC Special Research Report", Indicators of Terrorist Activity: Stopping the Next Attack In the Planning Stages, (2004).
- [14] A. Sands, "Integrating Open Sources into Transnational Threat Assessments," in Jennifer E. Sims and Burton Gerber, Transforming U.S. Intelligence, Washington: Georgetown University Press, (2005), pp. 67-68.
- [15] M. S. Soldatos and M. Hauswirth, "Convergence of Utility Computing with the Internet-of-Things", International Workshop on Extending Seamlessly to the Internet of Things (esIoT), IMIS-2012 International Conference, (2012).
- [16] A. N. Steinberg, C. L. Bowman and F. E. White, "Revisions to the JDL Data Fusion Model", in Sensor Fusion: Architectures, Algorithms, and Applications, Proceedings of the SPIE, vol. 3719, (1999).
- [17] H. Sundmaeker, P. Guillemin, P. Friess, S. Woelfflé, "Vision and Challenges for Realising the Internet of Things", (2010) March, ISBN 978-92-79-15088-3, doi:10.2759/26127, © European Union, (2010).
- [18] O. Vermesan and P. Friess, "Internet of Things - Global Technological and Societal Trends", The River Publishers Series in Communications, (2011).
- [19] I. G. Smith, O. Vermesan, P. Friess and A. Furness, "The Internet of Things 2012 New Horizons", ISBN: 978-0-9553707-9-3, http://www.internet-of-things-research.eu/pdf/IERC_Cluster_Book_2012_WEB.pdf.
- [20] A. Widder, R. Ammon, P. Schaeffer and C. Wolff, "Identification of suspicious, unknown event patterns in an event cloud", In Proceedings of the 2007 inaugural international conference on Distributed event-based systems (DEBS '07). ACM, New York, NY, USA, (2007).
- [21] A. Xue, W. Wang and M. Zhang, "Terrorist Organization Behavior Prediction Algorithm Based on Context Subspace", Advanced Data Mining and Applications, Lecture Notes in Computer Science, vol. 7121, (2011), pp. 332-345.
- [22] M. D. Turner and D. M. Weingberg, "A Simple Ontology for the Analysis of Terrorist Attacks", (2011), <http://hdl.handle.net/1928/13714>.
- [23] A. Najgebauer, R. Antkiewicz, M. Chmielewski and R. Kasprzyk, "The prediction of terrorist threat on the basis of semantic association acquisition and complex network evolution", Journal of Telecommunications & Information Technology, vol. 2008, Issue 2, (2008), p. 14.
- [24] B. Aleman-Meza, C. Halaschek, S. S. Sahoo, "TRAKS: Terrorist Related Assessment using Knowledge Similarity", Department of Computer Science, University of Georgia.

Authors



Stavros Petris, he received a B. Sc. Degree in Computer Systems (T.E.I. of Piraeus, 2005), a M. Sc. Degree in Data Communication Systems (Brunel, 2007) and a M. Sc. degree in Information Technology and Telecommunications (A.I.T., 2008). His experience, in the last five years, is in the area of software development for telematics applications and systems both in industry and academia. He is currently involved in two EU funded projects, namely, OpenIoT and PROACTIVE. His research interests involve grid computing middleware, internet-of-things, sensor networks and big data.



Christos Georgoulis, he holds a B. Sc in Computer Information Systems from the American College of Greece, an M. Sc in Business, Innovation and Entrepreneurship and an M. Sc in Information Technology and Telecommunications from AIT (Athens Information Technology). He has 7 years' experience in Project Management and Software Engineering in the context of European Research and Structural programs. His field of expertise involves front-end and back-end web development, sensor middleware and computer vision. Currently he is involved in three EU funded projects, namely OpenIoT, PROACTIVE and Vital.



John Soldatos, he holds a B. Sc degree and a PhD degree (2000) both from the National University of Athens. Since 1996 he has had very active involvement in more than fifteen research projects in the areas of broadband networks, pervasive/cloud computing, and the internet-of-things. He is the initiator and co-founder of open source projects Aspire RFID (<http://wiki.aspire.ow2.org>) and OpenIoT (<https://github.com/OpenIoTOrg/openiot>). He has published more than 140 papers in international journals and conferences. Since 2003 he is with Athens Information Technology, where he is currently an Associate Professor. He has also been an Adjunct Professor at the Information Networking Institute (of the Carnegie Mellon University (2007-2010) and a Honorary Research Fellow of the School of Computing of University of Glasgow (2014-2015).



Ilaria Giordani, she received the Master Degree in Computer Science at the University of Milano-Bicocca in 2006 and in 2010 the Ph. D in Computer Science at the same university with the thesis: "Relational Clustering for knowledge discovery in life sciences". She's currently working at Consorzio Milano Ricerche as researcher responsible for the execution of research projects in ICT domain founded by the European Commission, National and Regional. Her research activities are focused on the application of data analysis and data integration techniques in

different domains: security and aerospace and mobility, environment and logistics.



Raul Sormani, he received the Master Degree in Computer Science (2011) at the University of Milano-Bicocca, with the thesis: “HEMSYS: A system for the structural monitoring of helicopters”. He’s currently a second year PhD Student at the University of Milano-Bicocca, with a thesis on “Event detection in ambient intelligence”. He’s currently working at Consorzio Milano Ricerche as junior researcher on the European PROACTIVE project (PRedictive reasOning and multi-source fusion empowering Anticipation of attacks and Terrorist actions In Urban Environments). His researcher activities are focused on the analysis and development of reasoning algorithms for terrorist intent inference.



Divna Djordjevic, she received a PhD degree in Electronic Engineering and Computer Science from Queen Mary University of London, in 2006. She also holds a Dipl.-Ing. Degree in Telecommunications from the University of Nis in Serbia From 2008-2011 she was a Senior Researcher in the Data Analytics group at Accenture Technology Labs. She is currently a Senior Researcher at Consorzio Milano Ricerche. Her research interests include applied machine learning, data mining, social media analytics, and information retrieval.

