



HAL
open science

Trustworthiness Evaluation of Multi-sensor Situation Recognition in Transit Surveillance Scenarios

Francesco Flammini, Stefano Marrone, Nicola Mazzocca, Alfio Pappalardo, Concetta Pragliola, Valeria Vittorini

► **To cite this version:**

Francesco Flammini, Stefano Marrone, Nicola Mazzocca, Alfio Pappalardo, Concetta Pragliola, et al.. Trustworthiness Evaluation of Multi-sensor Situation Recognition in Transit Surveillance Scenarios. 1st Cross-Domain Conference and Workshop on Availability, Reliability, and Security in Information Systems (CD-ARES), Sep 2013, Regensburg, Germany. pp.442-456. hal-01506553

HAL Id: hal-01506553

<https://inria.hal.science/hal-01506553>

Submitted on 12 Apr 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Trustworthiness Evaluation of Multi-Sensor Situation Recognition in Transit Surveillance Scenarios

Francesco Flammini¹, Stefano Marrone², Nicola Mazzocca³ Alfio Pappalardo¹,
Concetta Pragliola¹, Valeria Vittorini³

¹ Ansaldo STS, Innovation & Competitiveness Unit, Naples, Italy
{francesco.flammini, alfio.pappalardo, concetta.pragliola}@ansaldo-sts.com

² Seconda Università di Napoli, Dip. Di Matematica e Fisica, Caserta, Italy
stefano.marrone@unina2.it

³ Università “Federico II” di Napoli, Dip. di Ingegneria Elettrica e Tecnologie
dell’ Informazione, Naples, Italy
{nicola.mazzocca, valeria.vittorini}@unina.it

Abstract. Physical Security Information Management (PSIM) systems are a recent introduction in the surveillance of critical infrastructures, like those used for mass-transit. In those systems, different sensors are integrated as separate event detection devices, each of them generating independent alarms. In order to lower the rate of false alarms and provide greater situation awareness for surveillance operators, we have developed a framework – namely DETECT – for correlating information coming from multiple heterogeneous sensors. DETECT uses detection models based on (extended) Event Trees in order to generate higher level warnings when a known threat scenario is being detected. In this paper we extend DETECT by adopting probabilistic models for the evaluation of threat detection trustworthiness on reference scenarios. The approach also allows for a quantitative evaluation of model sensitivity to sensor faults. The results of a case-study in the transit system domain demonstrate the increase of trust one could expect when using scenarios characterized in a probabilistic way for the threat detection instead of single-sensor alarms. Furthermore, we show how a model analysis can serve at design time to support decisions about the type and redundancy of detectors.

Keywords: Physical Security, Sensor and Data Analysis, Event Correlation, Trustworthiness, Probabilistic Modelling, Quantitative Evaluation

1. Introduction

In modern society the assurance of a secure environment is paramount due to the increasing number of threats against critical infrastructures. The number and the diversity of sensors used in modern wide-area surveillance is continuously increasing [1]. The type of sensors includes: (1) Environmental probes measuring temperature, humidity, light, smoke, pressure and acceleration; (2) Intrusion sensors, like magnetic

contacts, infrared/microwave/ultrasound motion detectors, etc.; (3) Radio-Frequency Identifiers (RFID) and position detectors, via satellite and/or electronic compasses; (4) Smart-cameras and microphones with advanced audio-video analytics capabilities; (5) Chemical Biological Radiological Nuclear explosive (CBRNe) detectors. Different types of sensing units are often integrated in smart-sensors like the so called 'motes' of Wireless Sensor Networks (WSN), featuring on-board 'intelligence' through programmable embedded devices with dedicated operating systems, processors and memory [5].

Such a wide range of sensors provides a large quantity of heterogeneous information which has to be handled properly, in terms of pre-processing, integration and reasoning, in order to effectively support PSIM operators; otherwise, there is the serious risk of overwhelming operators with unnecessary information, warnings or alarms, with the consequence of making them unable to perform their task and possibly underestimate critical situations [3][4].

In such a context, the issue of automatic situation recognition in PSIM is of paramount importance. However, not much work has been done in the research literature to develop frameworks and tools aiding surveillance operators to take advantage of recent developments in sensor technology. In other words, so far researchers seem to ignore the apparent paradox according to which the more and complex the sensors, the more and complex the tasks required for operators to manage and verify their alarms.

We have addressed the issue of automatic situation recognition by developing a framework for model-based event correlation in infrastructure surveillance. The framework – named DETECT – is able to store in its knowledge base any number of threat scenarios described in the form of Event Trees, and then recognize those scenarios in real-time, providing early warnings to PSIM users [6][7].

In this paper we adopt a model-based evaluation approach to quantitatively assess the effectiveness of DETECT in reducing the number of false alarms, thus increasing the overall trustworthiness of the surveillance system. The evaluation is dependent on sensor technologies and scenario descriptions, and it is based on stochastic modelling techniques. To achieve such an objective, some mappings are performed from Event Trees to other formalisms like Fault Trees, Bayesian Networks and Petri Nets (and their extensions). Those formalisms are widespread in dependability modelling and allow engineers to perform several useful analyses, including 'what if' and 'sensitivity', accounting for false alarms and even sensor hardware faults.

Generally speaking, the method used for the analysis, which is the main original contribution of this paper, allows to:

- Support design choices in terms of type and reliability of detectors, redundancy configurations, scenario descriptions.
- Demonstrate the effectiveness of the overall approach in practical surveillance scenarios, in terms of the increase of trustworthiness in threat detection with respect to single sensors.

In order to demonstrate the application of the methodology, a threat scenario of a terrorist attack in a metro railway station is considered.

The rest of this paper is structured as follows. Section 2 provides an overview of the related literature on DETECT and for trustworthiness evaluation of surveillance systems and it introduces the basic concepts of the event description language. Section 3 describes the process used for the analysis customizing it to the Bayesian Networks formalism in Section 4. Section 5 presents the case-study application using a metro-railway threat scenario. Finally, Section 6 provides the conclusions and hints for future improvements.

2 Background

The first concept of DETECT has been described in [6], where the overall architecture of the framework is presented, including the composite event specification language (EDL, Event Description Language), the modules for the management of detection models and the scenario repository. In [7], an overall system including a middleware for the integration of heterogeneous sensor networks is described and applied to railway surveillance case-studies. Reference [14] discusses the integration of DETECT in the PSIM system developed by AnsaldoSTS, namely RailSentry [2], presenting the reference scenario which will be also used in this paper. In order to detect redundancies while updating the scenario repository (off-line issue) and to increase the robustness of DETECT with respect to imperfect modelling and/or missed detections (on-line issue), distance metrics between Event Trees are introduced in [15].

A survey of state-of-the-art in physical security technologies and advanced surveillance paradigms, including a section on PSIM systems, is provided in [16]. Contemporary remote surveillance systems for public safety are also discussed in [17]. Technology and market-oriented considerations on PSIM can be also found in [18] and [21].

In [8] the authors address the issue of providing fault-tolerant solutions for WSN, using event specification languages and voting schemes; however, no model-based performance evaluation approach is provided. A similar issue is addressed in [9], where the discussion focuses on different levels of information/decision fusion on WSN event detection using appropriate classifiers and reaching a consensus among them in order to enhance trustworthiness. Reference [13] describes a method for evaluating the reliability of WSN using the Fault Tree modelling formalism, but the analysis is limited to hardware faults (quantified by the Mean Time Between Failures, MTBF) and homogenous devices (i.e. the WSN nodes). Performance evaluation aspects of distributed heterogeneous surveillance systems are instead addressed in [11], which only lists the general issues and some pointers to the related literature. Reference [10] about the trustworthiness analysis of sensor networks in cyber-physical system is apparently one of the most related to the topics of this paper, since it focuses on the reduction of false alarms by clustering sensors according to their

locations and by building appropriate object-alarm graphs; however, the approach is quite different from the one of DETECT and furthermore it applies to homogenous detectors. Another general discussion on the importance of the evaluation of performance metrics and human factors in distributed surveillance systems can be found in [12]; however, no hints are provided in that paper about how to perform such an evaluation on real systems.

Regarding the dependability modelling approach used in this paper, it is based on the results of the comparison among formalisms (i.e. Fault Trees, Bayesian Networks and Stochastic Petri Nets) in terms of modelling power and solving efficiency that has been reported in [20] and also applied in [19] to a different case-study using an approach known as ‘multi-formalism’.

2.1 Event Description Language

Threat scenarios are described in DETECT using a specific Event Description Language (EDL) and stored in a Scenario Repository. In this way we are able to permanently store all scenario features in an interoperable format (i.e. XML). A high level architecture of the framework is depicted in Fig. 1.




Fig. 1: The DETECT framework

A threat scenario expressed by EDL consists of a set of basic events detected by the sensing devices. An event is a happening that occurs at some locations and at some points in time. In this context, events are related to sensor data (i.e. temperature higher than a threshold). Events are classified as *primitive events* and *composite events*. A primitive event is a condition on a specific sensor which is associated with some parameters (i.e. event identifier, time of occurrence, etc...). A composite event is a combination of primitive events by means of proper operators. Each event is denoted by an *event expression*, whose complexity grows with the number of involved events. Given the expressions E_1, E_2, \dots, E_n , every application on them through any operator is still an expression. Event expressions are represented by *Event Trees*, where primitive events are at the leaves and internal nodes represent EDL operators.

DETECT is able to support the composition of complex events in EDL through a *Scenario GUI* (Graphical User Interface), used to draw threat scenarios by means of a user-friendly interface. Furthermore, in the operational phase, a model manager macro-module has the responsibility of performing queries on the Event History database for the real-time feeding of detection models corresponding to threat

scenarios, according to predetermined policies. Those policies, namely *parameter contexts*, are used to set a specific consumption mode of the occurrences of the events collected in the database. The EDL is based on the Snoop event algebra [24], considering the following operators: OR, AND, ANY, SEQ. For sake of space and due to their simplicity, the operators are not presented and further details are present in the literature.

3 Trustworthiness modelling process

The advantage of the modelling and analysis activity is twofold. On one hand it can be used during the design phase since it allows to quantitatively evaluate different design options for sensing and decision mechanisms allowing cost/effective trade-offs in protection systems design. In fact, the sensing strategies can differ in the number of sensors, in their reliability and/or in their efficiency in event detection; decision options are related to the logics that can be applied for correlating primitive events. On the other hand, the model can be used at run-time due to the possibility of tuning the models using data collected during the operational phase (i.e. event history log files merged with operator feedback about false negative/positive), allowing incremental refinement of detection models.

Fig. 2 shows how the aforementioned objectives can be achieved in an integrated process, in which both the monitored and monitoring systems are represented using probabilistic modelling formalisms. Quantitative model evaluation enables two possibilities:

- When used at design-time, the analyses can be used to compute the probability of having an alarm and its confusion matrix (i.e. the false positive and false negative probabilities). Such information can be used in order to improve the system by using more accurate or redundant sensors.
- When used at run-time, the detected events can be used as the evidence in the models. In such a way, the probability that the configuration of the primitive events is actually representative of the composite event (i.e. the threat scenario) can be dynamically adapted. Consequently, alarms can be generated only when the confidence in the detection is greater than a certain threshold.

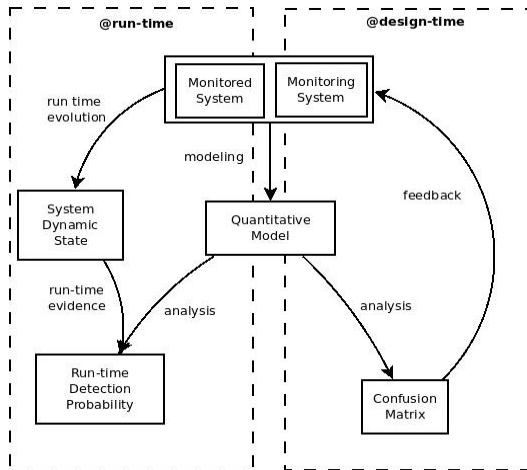


Fig. 2: The modelling and analysis process.




Fig. 3: Surveillance model layers.

Focusing on the design-time analysis, it is essential to develop an appropriate modelling methodology. In the context of surveillance systems trustworthiness evaluation, models of interest can be structured in three layers as depicted in

Fig. 3. These layers are:

- **Event layer:** this layer is devoted to modelling the actual cause-consequence relations in real environments. It determines how complex situations can be broken down into basic events (e.g. sneaking into a room by the window implies the breaking of the glass). It is usually the output of physical security surveys and risk assessment. In its most trivial form, it is constituted by the sequence of basic events associated to a threat scenario.
- **Sensing layer:** this layer models the sensors as objects with their characteristics (e.g. event detection capabilities, hardware reliability, detection performance) and the basic sensing actions with respect to the events identified in the lower layer.
- **Decision layer:** this layer addresses the (probabilistic) combination of simple events by means of EDL operators. It is important to note that this layer is built on top of the Sensing layer instead of the Event layer, since it does not deal with events actually occurring in the reality but with the ones generated by the sensing system, which can be different according to sensor types, deployment granularity, and detection performance.

In the following of this paper we mainly concentrate on the upper layer; however, the outputs of detection model evaluation can be used as inputs to refine threat modelling and better define sensor design parameters in order to meet the requirements of specific applications.

4 Application of the modelling process

In this Section we instantiate the process schema shown in Fig. 2 using the Bayesian Networks (BN) reference formalism, which features several advantages when employed in situation recognition. Fault Tree (FT) and Petri Net (PN) based processes can be equally derived from the general process schema. A complete comparison of these formalisms against their modelling power and efficiency is reported in [21]. In brief, FTs are very easy to build and analyse, but they have a limited modelling power. On the other hand, PNs feature a great expressive power but they are limited by the well-known state space explosion problem. BNs represent a good trade-off between those two extremes.

The operators used to build the Event Trees according to the event correlation approach implemented by DETECT have been briefly described in Section 2.2. Bayesian Networks fit the need to extend decision mechanisms adding the capability to handle probabilistic aspects. In fact features such as sensor hardware reliability and detection performance (i.e. false positive and false negative probabilities) rather than uncertainty in event modelling can be dealt with by appropriate BN subnets.

The process presented in Section 3 can be customized in the case of the BN formalism considering the specific types of analysis that can be conducted on a BN model [23]. Let: A be the set of alarms associated with threat scenarios; E be the set of events that can occur in the real environment; S be the set of states of the sensors. If we suppose $a \in A$, $e \in E$, $s \in S$, these three different indexes can be computed by solving the BN model:

- *Prior probability*, $P(a)$, that is the likelihood of occurrence of an alarm before any evidence relevant to the alarm has been observed. This index is the probability that an alarm is raised and it may be used at the design time of a PSIM system to predict the expected alarm rate, provided that the rate of primitive events is known a-priori.
- *Posterior probability*, $P(a | e, s)$, that is the conditional probability that an alarm is raised after some evidence is given. This index represents the probability of having an alarm in specific conditions, e.g. when some events happen (e.g. intrusion) and some others are generated by the surveillance system (e.g. sensor failure). It is useful at both design and run times. When used at design time it can be used to evaluate the performance of the detection system (i.e. the confusion matrix¹). In addition, the Posterior probability may be used to perform a ‘what-if’ analysis in order to evaluate the performance degradation in case of sensor failures. When used at run-time, a posterior analysis on the model fed with real evidence of events and/or sensor failures may provide a surveillance operator with alerts if probabilities are higher than a certain threshold.

¹ In this case of event detection, the confusion matrix accounts for binary events which can be *true* (i.e. occurred) or *false*. In DETECT, the false positive probability is given by $P(a = \text{true} | e = \text{false})$ while the false negative probability is $P(a = \text{false} | e = \text{true})$.

- *Likelihood*, $P(e | a, s)$, that is the probability of observing an element of E (real threat scenario) given evidence in A and S . In practice, it can be used to determine the probability that the alarm is trustworthy given that it has been generated. This kind of analysis is useful at run-time since it can support the decision making of the operators.

The layered model presented in Section 3 is substituted by a Bayesian Network where the BN nodes modelling the elements of the Event Layer are at the bottom, the ones representing the Sensing Layer are in the middle, the ones translating EDL operators at the Detection Layer are on the top. Specifically, we focus on the definition of a BN pattern that models the Sensing Layer and on the translation of EDL operators into BN elements. The BN pattern depicted in Fig. 4 shows how sensing can be modelled by means of three variables:

- **ev**: binary independent variable that models the occurrence of primitive events. The possible values of the variable are $\{true, false\}$;
- **sens**: binary independent variable representing sensor operation that can be $\{ok, down\}$;
- **det**: binary variable modelling event detection by the sensor. This is a $\{true, false, unknown\}$ dependent variable whose Conditional Probability Table (CPT) is reported in Tab. 1. The CPT is built considering the two probabilities: false-positive (sfp) and false-negative (sfn).

All the elements on the left side of Tab. 1 are translated into BN variables. The diverse operators are differentiated by CPTs.




Fig. 4: BN pattern for the Sensing Layer

Please note that, as combinatorial formalisms, Fault Trees and Bayesian Networks cannot precisely model the SEQ operator since they do not allow taking into account state and time dependant properties. In order to overcome such a limitation, more powerful formalisms are needed, like Dynamic Bayesian Networks or Petri Nets. However, it is possible to approximate a SEQ operator by an AND. In fact, since the SEQ requires the occurrence of events in a certain order, the set of cases in which e.g. SEQ ($E1, E2$) is true is a subset of the set in which AND ($E1, E2$) is true. Thus, by

substituting the SEQ with AND in the trustworthiness model, we are overestimating the false positive rate for the specific scenario.

Tab. 1: CPT of the Sensing Layer pattern

c	sens	det		
		true	false	unknown
false	down	0	0	1
false	ok	sfp	1-sfp	0
true	down	0	0	1
true	ok	1-sfn	sfn	0

5 Modelling trustworthiness in a specific scenario

The effectiveness of the modelling approach, described in the previous section, is demonstrated using a case-study in the mass transit domain, whose assets are vulnerable to several threats, including terrorist attacks. Therefore, surveillance systems for mass transit feature a growing number of heterogeneous sensing devices. In such a context, the quantitative evaluation of model trustworthiness and sensitivity to sensor faults is very important to design robust surveillance systems and to reduce the number of unnecessary alerts. In particular, at design time the results of model analysis provide valuable information to assess the level of redundancy and diversity required for the sensors, in order to find the appropriate configuration to comply with performance targets, perhaps given by the requirements specification of the end-user. Feedbacks from model evaluation can suggest changes about sensor dislocation and technologies. An estimation of detection model trustworthiness is essential also in real-time, whether using statically or dynamically updated data, in order to define confidence thresholds for triggering high level warnings and even automatic response actions.

Let us consider a threat scenario similar to the chemical attack with Sarin agent occurred in the Tokyo subway on March 20, 1995, which caused 12 fatalities and 5500 injured [22]. The available technologies to early detect and assess the threat include intelligent cameras, audio sensors and specific standoff CWA (Chemical Warfare Agents) detectors, which feature a limited alarm trustworthiness. By means of the DETECT framework, the events detected by these sensors could be correlated as well as reported in the threat scenario representation in the reference [14]. The main CWA detection technologies include Ion Mobility Spectroscopy (IMS), Surface Acoustic Wave (SAW), Infrared Radiation (IR), etc. They are employed in ad-hoc standoff detectors, characterized by different performances. One of the most accurate devices, the automatic passive IR sensor, can recognize a vapor cloud from several kilometres with a 87% detection rate. Obviously, it is possible to combine heterogeneous detectors (e.g. IMS/SAW and IR) and to correlate their alarms

according to different criteria (e.g. logic, temporal, and spatial), in order to increase the CWA detection reliability. The same considerations apply to the alarms detected by the other sensing devices.

The threat scenario consists of a simultaneous drop of CWA in subway platforms. Let us assume the following likely set of events:

1. attackers stay on the platforms, ready to drop the CWA;
2. contaminated persons fall down on the floor;
3. people around the contaminated area run away and/or scream;
4. CWA spreads in the platform level and possibly reaches higher levels.

In each subway site, it is possible to use two smart-cameras positioned at platform end walls, a microphone in the middle and two CWA standoff detectors positioned on the platform and on the escalators. The scenario can be formally described by means of the notation “sensor description (sensor ID) :: event description (event ID)”:

- *Intelligent Camera (S1) :: Fall of person (E1)*
- *Intelligent Camera (S1) :: Abnormal running (E2)*
- *Intelligent Camera (S2) :: Fall of person (E1)*
- *Intelligent Camera (S2) :: Abnormal running (E2)*
- *Audio sensor (S3) :: Scream (E3)*
- *IMS/SAW detector (S4) :: CWA detection (E4)*
- *IR detector (S5) :: CWA detection (E4)*

The Event Tree model of the CWA threat scenario is depicted in Fig. 5.

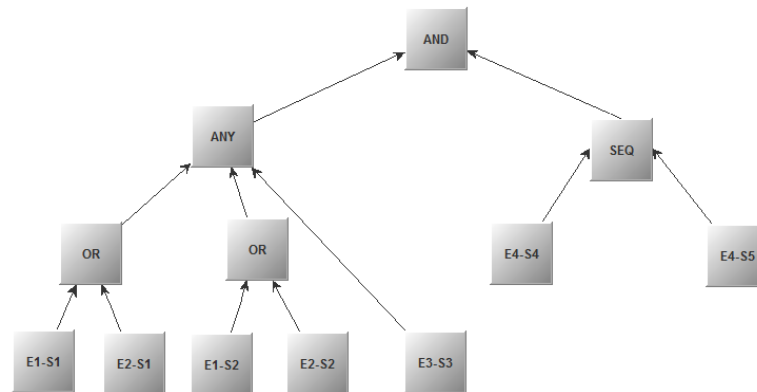


Fig. 5: Event tree associated to the CWA threat scenario

The OR operators correlate the events “person falling” and “person running”, detectable by the two redundant intelligent cameras monitoring the platform. The other child node (E3-S3) of the ANY operator represents the event “person screaming”, detectable by the intelligent microphone. When 2 out of these 3 events are detected in a certain (limited) time frame, the situation can reasonably be considered abnormal, so that a warning to the operator can be issued. The SEQ operator represents the upward CWA spread, detectable by the two redundant CWA sensors, installed at different levels. Finally the AND operator at the top of the tree represents the composite event associated with the whole CWA threat scenario.

As described in the previous section, each occurrence of an event can be *true* with a probability p , or *false* with a probability $1-p$. Each sensor can be available, i.e. *ok*, with a probability q , or unavailable, i.e. *down*, with a probability $1-q$. Finally, each single event detected by a sensor can be *true*, *false*, or *unknown* according to the occurrence of the event condition and to the availability of the sensor at that time. Moreover, each sensor, for each detectable event, is characterized by the values: sfn and sfp , which are the sensor false positive and false negative probabilities. The BN model of the event tree is built and analysed according to the modelling schema and methodology described in the previous sections and it is represented in Fig. 6.

The Event Layer is constituted by a node E that represents the actual CWA attack, while $E1$, $E2$, $E3$ and $E4$ are the primitive events that can be detected by the sensors. The interface between Event Layer and Sensor Layer is the set of $E1$, $E2$, $E3$ and $E4$ nodes. In the Sensor Layer, there are five nodes ($S1$, $S2$, $S3$, $S4$ and $S5$) representing sensors and seven nodes ($E1-S1$, $E2-S1$, $E1-S2$, $E2-S2$, $E3-S3$, $E4-S4$ and $E4-S5$) representing the sensed events. These seven event nodes constitute the interface between the Sensing Layer and the Detection Layers. Such an interface is built according to the mapping between EDL operators and BNs. As already stated, the SEQ operator has been substituted by the AND operator, introducing a modelling error.




Fig. 6: BN model of the CWA threat scenario

The model has been evaluated on the basis of the parameters summarized in Tab. 2, where (non conditional) probabilities refer to a standard time frame of 1 hour. The parameters have been valued considering realistic pseudo-data, since exact values depend on risk assessment results, specific sensor technology as well as operational reports in the real environment.

For the sake of brevity, we report only the posterior probability analysis that has been performed in order to evaluate the confusion matrix (see Tab. 3). The left column represents the evidence, that can be *true* (CWA threat is actually happening) or *false*. The other columns represent the probability of CWA threat alarm is generated ('Alarm on', which can be a true positive, *tp*, or false positive, *fp*, depending whether the evidence is true or false, respectively) or not ('Alarm off', which can be a *tn* or a *fn*, depending whether the evidence is false or true, respectively), or being inactive due to the unavailability of essential sensors. The results show that the rate of alarms, and in particular the *fp* and *fn* probabilities, is largely acceptable, according to recent ergonomics studies [4]. Furthermore, the value of *fp* is much less than false positives generated by single sensors. The evaluation of those parameters is essential to ensure system effectiveness and usability in real environments.

Tab. 2: BN model parameters

Name	Description	Node	Value
<i>attackProb</i>	Probability of having a CWA attack	E	10^{-6}
<i>running</i>	Probability of a running man in normal conditions (not related to an attack)	E1	$4*10^{-1}$
<i>falling</i>	Probability of a falling man in normal conditions (not related to an attack)	E2	10^{-3}
<i>screaming</i>	Probability of a scream in normal conditions (not related to an attack)	E3	$5*10^{-3}$
U_1	Unavailability of sensor 1	S1	$2*10^{-4}$
U_2	Unavailability of sensor 2	S2	$2*10^{-4}$
U_3	Unavailability of sensor 3	S3	10^{-4}
U_4	Unavailability of sensor 4	S4	$2*10^{-5}$
U_5	Unavailability of sensor 5	S5	10^{-5}
Sfp_{11} Sfp_{12}	Sensor false positive probability of sensor 1 (resp. 2) when sensing event 1	E1-S1	$3*10^{-2}$
Sfn_{11} Sfn_{12}	Sensor false negative probability of sensor 2 (resp. 2) when sensing event 1	E1-S2	$2*10^{-2}$
Sfp_{21} Sfp_{22}	Sensor false positive probability of sensor 1 (resp. 2) when sensing event 2	E2-S1	$2*10^{-2}$
Sfn_{21} Sfn_{22}	Sensor false negative probability of sensor 2 (resp. 2) when sensing event 2	E2-S2	$3*10^{-2}$
Sfp_{33}	Sensor false positive probability of sensor 3 when sensing event 3	E3-S3	$2*10^{-2}$
Sfn_{33}	Sensor false negative probability of sensor 3 when sensing event 3		$1.2*10^{-2}$
Sfp_{44}	Sensor false positive probability of sensor 4 when sensing event 4	E4-S4	$0.8*10^{-2}$
Sfn_{44}	Sensor false negative probability of sensor 4 when sensing event 4		$0.2*10^{-2}$
Sfp_{55}	Sensor false positive probability of sensor 5 when sensing event 5	E5-S5	$0.7*10^{-2}$
Sfn_{55}	Sensor false negative probability of sensor 5 when sensing event 5		$0.3*10^{-2}$

Tab. 3: Confusion matrix of the CWA threat scenario.

Evidence	Alarm on	Alarm off
True	0.995 (<i>tp</i>)	$0.22*10^{-4}$ (<i>fn</i>)
False	$0.5*10^{-2}$ (<i>fp</i>)	0.999978 (<i>tn</i>)

6 Conclusions and future work

Trustworthiness evaluation of models employed in situation assessment has a great practical importance in several applications of critical infrastructure surveillance. In those domains, quantitative evaluation is essential since the output of detection models is used to support decisions of the operators. Trustworthiness models allow to evaluate the robustness of PSIM systems also with respect to human errors and/or sensor faults, and to demonstrate compliance to performance and ergonomic requirements. In this paper, we have provided a structured trustworthiness modelling approach especially suited to surveillance systems featuring situation recognition capabilities based on Event Trees, which is the threat specification formalism used in the DETECT framework.

The effectiveness of the approach described in this paper is twofold. At design time, the results of the analysis provide a guide to support the choice and dislocation of sensors with respect to specific threats. At run-time, trustworthiness indices can be associated with detection models and hence to alarms reported to the operators, taking into account sensor performance and dependability parameters. Furthermore, at run-time:

- Sensor status (e.g. events detected, hardware failures, etc.) can be used to update trustworthiness indices in real-time
- The feedback of the operators over a significant time period can be used to fine-tune trustworthiness parameters (e.g. the *fp* probability can be estimated by counting the average number of false alerts generated by single sensors or even by DETECT, and by normalizing that number according to the reference time frame).

We have shown that among the probabilistic modelling formalisms, BNs are the most suited to this kind of application, allowing a very good trade-off between ease of modelling and expressive power.

The results achieved by model evaluation demonstrate the effectiveness of the DETECT event correlation approach to reduce the number of unnecessary alerts, warning and alarms, thus improving PSIM ergonomics and usability. Model evaluation also allows to perform 'what-if' predictions and sensitivity analyses with respect to changes in detection model structure and parameters, enabling and supporting design optimisation at several levels.

Future developments will address the following: evaluation results are going to be extended using further models and simulation campaigns; data coming from on-the-field experimentations and long term observations is going to be integrated with the models and used to validate them. The aforementioned automatic update of trustworthiness parameters is being implemented in DETECT using appropriate modules and exploiting the integration of DETECT in the PSIM system developed by AnsaldoSTS.

References

1. Garcia, M.L.: The Design and Evaluation of Physical Protection Systems. Butterworth-Heinemann (2001)
2. Bocchetti, G., Flammini, F., Pragliola, C., Pappalardo, A.: Dependable integrated surveillance systems for the physical security of metro railways. In: IEEE Procs. of the third ACM/IEEE International Conference on Distributed Smart Cameras (ICDSC 2009), pp. 1--7 (2009)
3. Zhu, Z., Huang, T.S.: Multimodal Surveillance: Sensors, Algorithms and Systems. Artech House Publisher (2007)
4. Wickens, C., Dixon, S.: The benefits of imperfect diagnostic automation: a synthesis of the literature. In: Theoretical Issues in Ergonomics Science, 8(3), pp. 201--212 (2007)
5. Flammini, F., Gaglione, A., Mazzocca, N., Moscato, V., Pragliola, C.: Wireless Sensor Data Fusion for Critical Infrastructure Security. In: Advances in Intelligent and Soft Computing, Volume 53, pp. 92--99 (2009)
6. Flammini, F., Gaglione, A., Mazzocca, N., Pragliola, C.: DETECT: a novel framework for the detection of attacks to critical infrastructures. In: Safety, Reliability and Risk Analysis: Theory, Methods and Applications, Martorell et al. (Eds), Procs of ESREL'08, pp. 105--112 (2008)
7. Flammini, F., Gaglione, A., Mazzocca, N., Moscato, V., Pragliola, C.: On-line integration and reasoning of multi-sensor data to enhance infrastructure surveillance. In: Journal of Information Assurance and Security (JIAS), Vol. 4, Issue 2, pp. 183--191 (2009)
8. Ortmann, S., Langendoerfer, P: Enhancing reliability of sensor networks by fine tuning their event observation behavior. In: Proc. 2008 International Symposium on a World of Wireless, Mobile and Multimedia Networks (WOWMOM '08). IEEE Computer Society, Washington, DC, USA, pp. 1-6
9. Bahrepour, M. and Meratnia, N. and Havinga, P.J.M.: Sensor Fusion-based Event Detection in Wireless Sensor Networks. In: 6th Annual International Conference on Mobile and Ubiquitous Systems: Networking and Services, MobiQuitous 2009, 13-16 July 2009, Toronto, Canada.
10. Lu-An Tang, Xiao Yu, Sangkyum Kim, Jiawei Han, Chih-Chieh Hung, and Wen-Chih Peng. 2010. Tru-Alarm: Trustworthiness Analysis of Sensor Networks in Cyber-Physical Systems. In Proceedings of the 2010 IEEE International Conference on Data Mining (ICDM '10). IEEE Computer Society, Washington, DC, USA
11. Legg, J. A.: Distributed Multisensor Fusion System Specification and Evaluation Issues. Defence Science and Technology Organisation, Edinburgh, South Australia 5111, Australia, October, 2005
12. Karimaa A.: Efficient Video Surveillance: Performance Evaluation in Distributed Video Surveillance Systems. In: Video Surveillance, Weiyao Lin (Ed.), 2011, ISBN: 978-953-307-436-8, InTech.
13. Silva I., Guedes L.A., Portugal P., Vasques F.: Reliability and Availability Evaluation of Wireless Sensor Networks for Industrial Applications. In: Sensors. Vol. 12, No. 1 2012: pp. 806-838
14. Flammini, F., Mazzocca, N., Pappalardo, A., Pragliola, C., Vittorini, V.: Augmenting surveillance system capabilities by exploiting event correlation and distributed attack

- detection. In: Proc. 2011 Intl. Workshop on Security and Cognitive Informatics for Homeland Defence (SeCIHD'11), co-located with ARES'11, A. M. Tjoa et al. (Eds.), LNCS 6908, pp. 191-204
15. Flammini, F., Pappalardo, A., Pragliola, C., Vittorini, V.: A robust approach for on-line and off-line threat detection based on event tree similarity analysis. In: Proc. Workshop on Multimedia Systems for Surveillance (MMSS) in conjunction with 8th IEEE International Conference on Advanced Video and Signal-Based Surveillance, Klagenfurt, Austria, August 29-30, 2011: pp. 414-419
 16. Flammini, F., Pappalardo, A., Vittorini, V.: Challenges and emerging paradigms for augmented surveillance. In: Effective Surveillance for Homeland Security: Combining Technology and Social Issues. Taylor & Francis / CRC Press, 2013. To appear.
 17. Rätty, T. D.: Survey on contemporary remote surveillance systems for public safety. In: IEEE Trans. Sys. Man Cyber Part C, No. 40, Vol. 5, Sep. 2010: pp. 493-515
 18. Hunt, S.: Physical security information management (PSIM): The basics. Available at: <http://www.csoonline.com/article/622321/physical-security-information-management-psim-the-basics>
 19. Flammini, F., Marrone, S., Mazzocca, N., Vittorini, V.: A new modelling approach to the safety evaluation of N-modular redundant computer systems in presence of imperfect maintenance. In: Reliability Engineering & System Safety, Vol. 94, Issue 9, September 2009: pp. 1422-1432
 20. Bobbio, A., Ciancamerla, E., Franceschinis, G., Gaeta, R., Minichino, M., Portinale, L.: Sequential application of heterogeneous models for the safety analysis of a control system: a case study. In: Reliability Engineering & System Safety journal, RESS, Vol. 81, No. 3, 2003: pp. 269-280
 21. Frost & Sullivan: Analysis of the Worldwide Physical Security Information Management Market. November 2010. Available at: http://www.cnsoftware.com/media/reports/Analysis_Worldwide_Physical_Security_Information_Management_Market.pdf
 22. National Consortium for the Study of Terrorism and Responses to Terrorism (START). (2012). Global Terrorism Database [199503200014]. Retrieved from <http://www.start.umd.edu/gtd>
 23. Charniak, E.: Bayesian Networks without Tears, AI Magazine, 1991
 24. Chakravarthy, S., Mishra, D.: Snoop, An expressive event specification language for active databases. Data Knowl. Eng., Vol. 14, No. 1, 1994: pp. 1-26
 25. D. Codetta-Raiteri, The Conversion of Dynamic Fault Trees to Stochastic Petri Nets, as a case of Graph Transformation, Electronic Notes in Theoretical Computer Science, vol. 127(2), pages 45-60, Elsevier, March 2005