Event Detection Using Abductive Reasoning on Sensor Data

Bill Karakostas
VLTN BBVA
Antwerp, Belgium
bill.karakostas@vltn.be

Abstract— **We present an approach that detects physical events such as a fire or an explosion using sensor data fusion, where not all relevant signals describing the event are available due to non-presence or malfunctioning of some sensors. We employ abductive probabilistic reasoning to detect the occurrence of an event amongst several alternative events from imperfect sensor data. Influenced by Dempster-Shafer's evidence theory, we reason on the available evidence produced by the sensor data, combined with counterevidence, to establish degrees of confidence to the different hypotheses made about the occurrence of an event. The paper also describes an experimental sensor setup for detection of fire and explosion events, and its effectiveness in terms of false negative and false positive detection rates.**

Keywords- sensor data fusion; event detection; abductive reasoning; Raspberry Pi.

I. INTRODUCTION

Sensor data fusion is the method of combining data from homogeneous or heterogeneous multiple sensors in order to form a unified picture [1]. Compared to data obtained from single sensors, multisensor data fusion improves the overall event detection capabilities, in terms of reduction in the false positive and false negative detection rates. Data fusion systems are now widely used in various areas such as sensor networks, robotics, video and image processing, and increasingly so, in Internet of Things (IoT) applications such as smart cities [2].

However, in sensor deployments, some detection capabilities may not be available because of lack of suitable sensors or because of the quality (precision, accuracy, reliability) of the obtained sensor data. Hence, false positives which means detection of a non-existing event, or false negatives which means failure to detect an event, may occur.

Therefore, intelligent processing of sensor data may be needed to rectify such deficiencies. Inference on the sensor data entails the ability to (a) detect that an event has occurred and (b) determine the type of this event amongst a number of possible event types. Logic based approaches have been employed for such purposes, however not all types of logical inferences are possible, due to incomplete data and/or weak causal relationships between an event and its manifestations. Because deductive inferencing of the event from its manifestations is not always justifiable, forms of inductive (probabilistic) [3] and abductive reasoning [4] have been employed. In this paper, we employ a variant of the later form of reasoning, which calculates evidence about the occurrence of an event from sensor data obtained from multiple sensors, and also counter-evidence from lack of observed data. Each type of sensor data is assigned a numerical weight to indicate the degree to which presence of such data supports the evidence about the occurrence of an event. Evidences and counter-evidences are compared using likelihood ratio methods, across the range of possible events, in order to find the event with the highest evidence ratio. Additionally, confidence in the suggested event detection is calculated in terms of sensitivity and specificity of the sensor layout used for event detection.

The structure of the paper is as follows. The next section presents related work, while the theoretical model for our sensor event detection method is introduced in Section III. Section IV discusses an experimental setup and experiment results for detecting and classifying smoke and explosion events from sensor data. This section also analyses the effectiveness of the approach. Section V discusses the advantages and limitations of the proposed approach to event detection and classification from sensor data, and proposes future research.

II. RELATED WORK

The approach described in this paper aims to provide a reliable way for detecting and classifying physical phenomena (events) from fused data collected from potentially unreliable measurements/sensors, to which different weights are assigned as evidence. By applying evidence-based theory, we attempt to calculate and compare the likelihoods of occurrence for the different types of events. To compensate for measurement/sensor unreliability we combine the evidence from multiple events, including the absence of evidence. Other approaches have also utilised Dempster Schafer theory for sensor data fusion [6]. Also, some approaches combine the Dempster- Shafer evidence theory with other machine learning techniques such as hierarchical neural networks, to improve the accuracy of classifications [7].

 ISBN: 978-1-61208-659-0

III. THEORETICAL MODEL

Physical events as opposed for example to events occurring in the digital domain, are characterised by physical processes and their quantities (energy, light, sound). For example, a fire is manifested usually by an increase in the ambient temperature and the presence of gasses. However, the manifestations of different instances of an event may vary. For example, although fires in general produce smoke, depending on the type of materials combusted, some fires may produce very little or no smoke. The intensity of the heat, the volume and composition of smoke and other physical characteristics are also subject to many parameters in the environment of the fire. Also, different types of events may have similar manifestations, for example both a fire and an explosion may produce smoke. Finally, there are detection (e.g. accuracy) limitations imposed by the technology used to manufacture the sensors, the sensor deployment layout, as well as by possible sensor malfunctioning. Such limitations can restrict our ability to use sensor data, by failing to detect an event (false negative), falsely detect an event that did not occur (false positive), or by wrongly classifying an event. Therefore, our model addresses the inherent uncertainty in event detection from sensor data and employs a probabilistic approach, influenced from concepts from Dempster-Shafer evidence theory [4] to reason on the available evidence produced by the sensors and combine it with counterevidence to establish degrees of confidence in the various hypotheses made about the occurrence of events.

It must be noted that evidence theory has been utilised for similar purposes such as diagnostic tasks where information is also obtained from sensors [5]. However, our aim is to obtain reliable information by fusing a mixture of both reliable and unreliable data and also from lack of data measurements.

Our approach is formally described as follows: Let E be the set of all types of phenomena we consider for detection and classification and M the set of all manifestations of events in E, detectable by our sensors, with each e in E characterised by a set of manifestations $M_e \subset M$. We define a weighting function w_e that for every m in M_e assigns a value in {0,1} to m. The weight produced by function $w_e(m)$ represents the degree to which observation of manifestation m increases the evidence that phenomenon e is occurring. We also define weights to measure *lack of evidence* i.e. the degree to which non-observation of a manifestation supports the evidence that the phenomenon has not occurred. Thus, w assigns a weight in {0,1} for each non-manifestation $\neg m$ as a measure of the evidence that the lack of a manifestation provides to support that the phenomenon has not occurred.

Note that the values for evidence and counterevidence do not have to be correlated. For example, the sensing of heat is a strong indicator that a fire has occurred, however the absence of heat detection is not an equally strong indicator that a fire has not occurred, as some fires initially do not produce measurable heat, or there is a possible malfunction of the heat sensing sensor.

Finally, we define the total evidence from our observations $O_e \subset M_e$ and also from the lack of them i.e. $\acute{O}_e = M_e - O_e$ about a phenomenon of type e with manifestations M_e as

$$Ev(e) = \sum_{m \in O_e} w(m) \quad (1)$$

$$Ev(\neg e) = \sum_{m \in \acute{O}_e} w(m) \quad (2)$$

The two formulas for evidence and counterevidence allow us to reason in a qualitative manner about the likelihood of a phenomenon as well as to compare the likelihoods of different phenomena. For example, even if the counterevidence for a phenomenon is zero, it is not justifiable to conclude that the phenomenon has occurred, if the gathered evidence for the phenomenon is also small. The ratio of evidence to counterevidence is an indicator of the degree of certainty from the observations. In particular, if Ev(e) ~ Ev(¬e) it indicates uncertainty as to whether the phenomenon has occurred or not. If the evidence weight Ev(e) is equal to the theoretical maximum $|M_e|$ i.e. $O_e = M$, then there is perfect evidence, as by necessity Ev(¬e) has to be zero. Correspondingly, total lack of evidence occurs when Ev(¬e)= $|M_e|$.

To further strengthen the confidence in the diagnosis of events, prior knowledge about the frequency of occurrence of different events, if available, can be utilised. This can be expressed as a probability value *P(e)* with values in {0,1}. Probabilities are used as prevalence values for calculating likelihood ratios as explained in the next section.

As different types of events may have similar manifestations, we are often interested in being able to determine which is the most likely event has occurred, i.e. to correctly classify events. We are also interested in not falsely identifying an event (false positive) and avoid failing to detect an event (false negative) For this purpose, we utilise likelihood ratios.

We define likelihood ratio (LR) as the ratio of the probability that the event has been correctly identified, to the probability that the event has been incorrectly identified. We use sensitivity and specificity of the measurements as the numbers used to generate a LR. Sensitivity is the proportion of truly occurred events that are identified as such by the system (i.e. their total evidence scores are the highest amongst all candidate events). Specificity is the proportion of non-events that have been correctly identified and indicate the probability that the test will correctly identify a non-event.

We calculate LR for both positive and negative event identifications, expressed as 'LR+' and 'LR-', respectively. The calculations are based on the following formulas:

LR+ = sensitivity / 1- specificity (3)

LR- = 1- sensitivity / specificity (4)

 ISBN: 978-1-61208-659-0

IV. EXPERIMENTATIONS

A. Experimental apparatus

We have empirically validated our approach in a sensor setup to detect and classify fire and explosion events. Fig. 1 shows the experimental device setup. On the left of the picture there is a Raspberry Pi Model B single board computer, with a GrovePi sensor HAT (‘hardware at top’) board, to which GrovePI sensors for atmospheric pressure, air quality, light intensity, sound and temperature are attached (some of the sensors are visible at the left of the picture). Small scale fires and explosions were produced under controlled conditions, in order to obtain the experimental data. Data were collected from the sensors periodically and stored by the Raspberry Pi where they are analysed for event detection. An event is detected, if the rate of increase in the values read by the sensor exceed a threshold, the rate of change is calculated using the five point formula $r_t = (v_{t-2} - 8v_{t-1} + 8v_{t+1} - v_{t+2})/12$, where v_t are data measurements at different timepoints. Fig. 2 shows examples of time series of collected sensor data that shows the manifestations of the different events (air pressure, temperature and light intensity), through the increase in rates in the recorded sensor values. The two events share some manifestations, for example heat, but no change to the air pressure. Additionally, the weights assigned to each manifestation differ per event. We assume that weights are assigned to manifestations by experts and that the specifics of the environment where the sensors are deployed are taken into account, to calibrate the values of the weights.

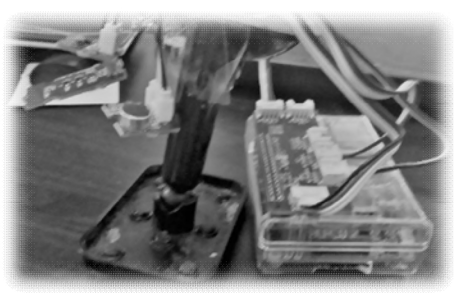

Figure 1. Apparatus used for experiments.

B. Assigning weights to phenomena

Table I demonstrates the above approach in the modelling of two physical events, a fire and an explosion. The two events share some manifestations such as heat but not others, for example, air pressure. Additionally, the weights assigned to each manifestation differ per event. We assume that weights are assigned to manifestations by experts and that the specifics of the environment where the sensors are deployed are taken into account.

C. Calculating Evidences

Table II shows the aggregated evidences for the two types of events, as well as for their negations (non-events). Additionally, the second column of Table II shows the manifestations detected by the sensors, and the third column of Table II shows the mechanisms used to generate fire and explosion manifestations.

The experiment results were used to evaluate the detection and classification capabilities of our system. In experiment #2, although the evidence of fire is overwhelming compared to counterevidence, the system cannot clearly make the case for or against an explosion with weights of 1.4 and 1.3, respectively. In experiment #3, the system narrowly suggests the hypothesis of fire over non-fire, while clearly rejects the hypothesis of an explosion. Evidence for and against fire are however (correctly) insufficiently high for the fire hypothesis to be strongly suggested.

D. Likelihood ratios, sensitivity and specificity calculations

The purpose of the experiments was to estimate the detection accuracy of our approach. Detection accuracy is the percentage of correct event detections when taking into account the total percentage of false positives and false negatives. A detection accuracy of 100% would imply a 0% false positives and false negatives rate, but may not necessarily imply zero classification errors, as explained below.

Classification errors refer to the number of misdiagnoses made during the experiment. To calculate the relevant likelihood ratios, prevalence of the explosion and fire events were set to 0.01 and 0.05 respectively, meaning that in general, 1% of detected events are true explosions and true fires respectively, also suggesting that fires are five times more common than explosions. From the data of Table II, specificities and sensitivities for fire and explosion were calculated, and then Formulas (3) and (4) were applied to calculate positive and negative likelihoods (LR+, LR-). The results are summarised in Table III.

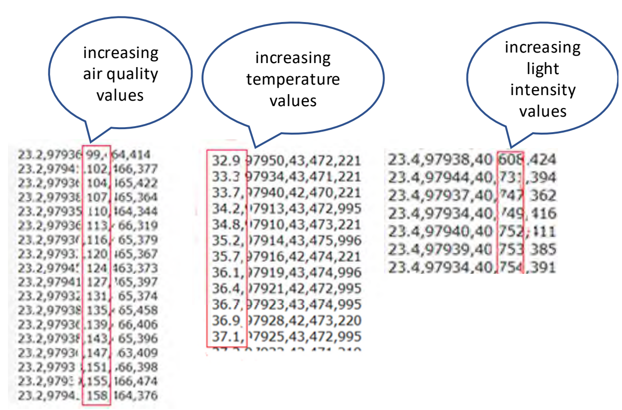

Figure 2. Time series of sensor data showing manifestation of phenomena. Highlighted areas from left to right show how values for: air quality (smoke), temperature and light intensity are increasing due to the event occurrence.

 ISBN: 978-1-61208-659-0

TABLE I. SENSITIVITIES, SPECIFICITIES AND LIKELIHOOD RATIOS (+/-) FOR FIRE AND EXPLOSION EVENTS

Event Type	*Sensitivity*	*Specificity*	*LR+*	*LR-*
Fire	1/4	0	1/4	n/a
Explosion	3/4	1/2	3/8	1/8

TABLE II. SENSITIVITIES, SPECIFICITIES AND LIKELIHOOD RATIOS (+/-) FOR FIRE AND EXPLOSION EVENTS

Event Type	*Sensitivity*	*Specificity*	*LR+*	*LR-*
Fire	1/4	0	1/4	n/a
Explosion	3/4	1/2	3/8	1/8

TABLE III. EVIDENCE AND COUNTEREVIDENCE FOR DIFFERENT EVENTS

Exp #	**Parameters: (F: fire evidence, E:explosion evidence)**					
	Manifestations	***Mechanism used***	***F***	***¬F***	***E***	***¬E***
1	Sudden pressure change (detonation) flash and loud noise, smoke, no heat	Firecracker	1.4	0.01	2.8	0.6
2	Flash of light, heat, smoke, no high pressure or loud noise	Ignition of flammable Liquid	1.3	0.03	1.4	1.3
3	Flash of light, no heat, sound, smoke or high pressure	Flashlight	0.2	0.14	0.8	2.0
4	Smoke, no heat, light or high pressure	Cigarette smoke	0.9	0.04	0.3	1.2
5	Heat with flame, some smoke, no noise or high pressure	Bunch of matches ignited	1.3	0.03	1.4	1.3
6	Explosion (no light or heat, Noise and high air pressure	Burst balloon	0.31	0.12	1.7	0.8
7	Heat, no high pressure, light or noise	Heatgun	0.5	0.13	1.2	1.4

V. DISCUSSION AND CONCLUSIONS

Our approach implicitly considers sensor reliability in assigning evidence weights to sensor data. An explicit modelling of sensor unreliability would allow additional scenario to be produced. Also, we use crisp thresholds for the detection of manifestations, i.e. a manifestation has either been detected or not. In actuality, the boundaries between a manifestation and a non-manifestation may be fuzzy rather than crisp. For example, the distinction between a bright light and a flash of light may be better defined in terms of membership to a fuzzy set of values. Since LR+ shows how much more likely it is for a truly occurred event to score higher than a non-event, Table III suggests that the above sensor setup is more suitable for detecting explosions rather than fires, while minimising false positive rate. This is both due to the relative assignment of weights to different phenomena and to the types of sensors used for detection. This could indicate that more fire specific detection sensors would need to be deployed to improve sensitivity and specificity.

Another limitation of our approach relates to the need to assign evidence weights to each manifestation per corresponding event. Such weights can be obtained empirically by historical data recorded over long periods (where for example multiple explosion and fire events occurred with their manifestations statistically analysed). Where such data are not available and/or the particular context characteristics must also be accommodated (e.g. the topology/layout of the area where the explosion or fire occurs), expert opinion is necessary,

Amongst the advantages of the proposed approach, we include the clarity of the model and the simplicity of the calculations it employs. This unlike for example, approaches based on neural networks, makes the model easily open to inspections and necessary calibrations. Additionally, since the calculations involved are rather rudimentary, devices of limited computational power and storage can be employed, such as gateways and other ‘edge’ devices that are physically deployed closer to the sensors.

For further research, our approach could include more sophisticated pattern matching techniques for detecting event manifestations, combination with other machine learning techniques such as neural networks, and integration with intelligent processing and decision support systems used for example, in risk management.

ACKNOWLEDGEMENT

Work described in this paper was financially supported by EU Horizon 2020 Project CHARIOT (Cognitive Heterogeneous Architecture for Industrial IoT- Grant No. 780075).

REFERENCES

[1] B. Khaleghi, A. Khamis, O. Fakhreddine, S. Karray and N. Razavi. “Multisensor data fusion: a review of the state-of-the-art”. Information Fusion. Volume 14 Issue 1, January, 2013 pp. 28-44.

[2] M. Wang, P. Charith, Jayaraman, P. Prem, M. Zhang, P. Strazdins, and R. Ranjan. “City data fusion: sensor data fusion in the internet of things”. International Journal of Distributed Systems and Technologies (IJDST), 2015.

[3] J. Pearl. Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference. CA: MKP, 1988.

[4] G. Shafer. A Mathematical Theory of Evidence. Princeton University Press, 1966.

[5] E. Pashaa, H.R. Mostafaeib, M. Khalajc and F. Khalajb “Fault diagnosis of engine using information fusion based on Dempster-Shafer theory”. J. Basic. Appl. Sci. Res., 2(2)1078-1085, 2012.

[6] H. Wu, M. Siegel, R. Stiefelhagen and Yang J. “Sensor fusion using Dempster-Shafer theory”. IEEE Instrumentation and Measurement Technology Conference Anchorage, AK, USA, 21-23 May 2002.

[7] R. Fay, F. Schwenker, C. Thiel and G. Palm. “Hierarchical neural networks utilising Dempster-Shafer evidence theory”. Artificial Neural Networks in Pattern Recognition: Second IAPR Workshop, ANNPR 2006, Ulm, Germany, August 31-September 2, 2006, pp.198-209.