




Large Scale Surveillance, Detection and Alerts Information Management System for Critical Infrastructure

Z. Sabeur¹ , Z. Zlatev¹, P. Melas¹, G. Veres¹, B. Arbab-Zavar¹, L. Middleton¹, and N. Museux²

¹ Department of Electronics and Computer Science, University of Southampton IT Innovation Centre, Southampton, UK

{zas, zdz, pm, gvv, baz, ljm}@it-innovation.soton.ac.uk

² THALES Research and Technology, Palaiseau, France

nicolas.museux@thalesgroup.com

Abstract. A proof-of-concept system for large scale surveillance, detection and alerts information management (SDAIM) is presented in this paper. Various aspects of building the SDAIM software system for large scale critical infrastructure monitoring and decision support are described. The work is currently developed in the large collaborative ZONeSEC project (www.zonesec.eu). ZONeSEC specializes in the monitoring of so-called Wide-zones. These are large critical infrastructure which require 24/7 monitoring for safety and security. It involves integrated in situ and remote sensing together with large scale stationary sensor networks, that are supported by cross-border communication. In ZONeSEC, the specific deployed sensors around the critical infrastructure may include: Accelerometers that are mounted on perimeter fences; Underground acoustic sensors; Optical, thermal and hyperspectral video cameras or radar systems mounted on strategic areas or on airborne UAVs for mission exploration. The SDAIM system design supports the ingestion of the various types of sensors platform wide-zones' environmental observations and provide large scale distributed data fusion and reasoning with near-real-time messaging and alerts for critical decision-support. On a functional level, the system design is founded on the JDL/DFIG (Joint Directors of Laboratories/Data Fusion Information Group) data and information fusion model. Further, it is technologically underpinned by proven Big Data technologies for distributed data storage and processing as well as on-demand access to intelligent data analytics modules. The SDAIM system development will be piloted and alidated at various selected ZONeSEC project wide-zones [1]. These include water, oil and transnational gas pipelines and motorway conveyed in six European countries.

Keywords: Big data · Data fusion · Information systems
Surveillance of critical infrastructure

1 Introduction

The SDAIM functional objectives specialize in the enablement of the intelligent fusion of data and information processing and reasoning from heterogeneous observation data sources. These are generated from a high variety of sensor observation platforms and processing components. Currently, the SDAIM considers observation data sources from 3D accelerometers, underground acoustic sensors, CCTV cameras, thermal and hyper-spectral cameras, radars, SCADA (Supervisory Control and Data Acquisition) systems and human observers. This is achieved by undertaking sources, data and information modelling through the creation of metadata for such sources, in order to automate the system and data processing configurations. The aim is to achieve sensors platforms “plug-and-play” and the automatic on-demand access to fusion processes, their configuration and execution. Furthermore, open standards are used for fetching sensor data and processing metadata such as OGC Sensor Web Enablement using SensorML, JSON and RDF; and OWL for metadata description and modelling.

In addition to the above, the SDAIM system encounters high big data velocities for processing under the fusion resources. This is the case because of high data transmission volumes from thousands of sensors mounted at various platforms. This challenging problem is addressed through combining of advanced system’s architecture design solutions with the adoption of specialized Big Data technologies. Specifically, the SDAIM advanced architecture is based on the de-facto functional data and information fusion model, the JDL/DFIG generic model [2–4]. This modelling framework provides highly structured functionalities of the SDAIM data processing modules, while adopting the current Google approach for scalable big data processing workloads. The latter uses concepts of Dockers for creating lightweight virtualized containers for the processing modules, and Kubernetes for scalable deployment and execution of processing components. This approach enables the SDAIM to be deployed and executed in a distributed fashion over a high variety of operating systems, ranging from lightweight embedded platforms to heavy-duty servers.

2 SDAIM Scalable Data Fusion Approaches

An important aspect with regards to the scaling of the SDAIM data fusion and processing is our structured approach for achieving a scalable high level data and information fusion. This particularly needs to be addressed at the JDL fusion level 3 on Situation Assessment. In the SDAIM, level 3 fusion components generate, what we call, ‘alerts’ and these needs to be critically achieved through an intelligent processing of big data for messaging them into a distributed system with effective action and mitigation on detected illicit activities at wide-zones. The JDL framework, see Fig. 1, is defined as follows¹:

¹ At this stage, only the first four levels of the JDL/DFIG framework are described. The full extended level will be described in the full paper as defined in Fig. 1.

Level 0 → Fusion components generate signals, e.g. raw signal, co-registered signals with early data pre-processing, harmonisation and aggregations
 Level 1 → Fusion components which specialise in the identification of background low level feature within a scene of interest
 Level 2 → Fusion components which detect events concerning objects of targeted types; and/or critical state transitions within a given sensory signal; and/or specific behaviours which may be qualified as unusual
 Level 3 → Fusion components which generate alerts on the detected objects, events or behaviour of critical status that may compromise the security and safety of the wide-zone of interest. These fusion components work on the level of reasoning on objects and behaviours relations.




Fig. 1. The JDL/DFIG data and information fusion model

Alerts are consequently disseminated for consideration for decision-support to the security practitioners’ operations via an online accessible graphical user interface. It will represent the Common Operational Picture (COP) of the wide-zone specific scene of interest. Depending on the COP diagnosis, the security practitioners could potentially raise a level of alert to a specified hierarchy of decision-makers and first responders.

Based on the JDL/DFIG information model, the logical architecture of the SDAIM is depicted on Fig. 2. In it, the central components are the data and information fusion algorithms. These components, as per the JDL/DFIG model, fulfil fusion functions at different logical levels, e.g. signal processing, event or object detection, multiple classifications with various levels of confidence to high level extracted knowledge through fusion and reasoning on the wide-zone operational spatial and temporal contexts.




Fig. 2. SDAIM architecture logical view

Finally, Fig. 3 depicts a simplified deployment case of the SDAIM. The SDAIM is deployed for a Wide-zone with 3 sub-regions and their respective sub-regional and global regional control centers.

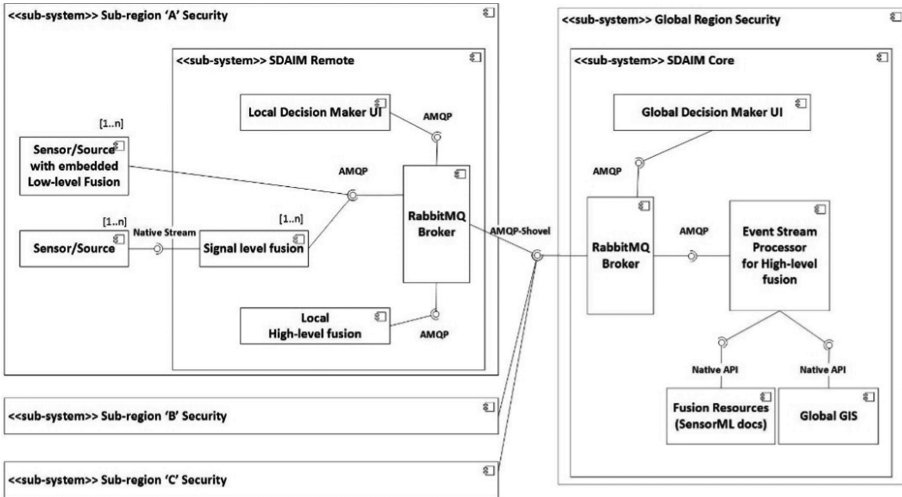


Fig. 3. SDAIM deployment for a wide-zone with 3 sub-regions and sub-regional and a global regional control centres

3 Applications on Surveillance, Detection and Alerts of Illicit Behaviour in Wide-Zones

In this section, we present early specific applications concerning surveillance, automated detection and alerts of illicit behaviour in wide-zones. The use of multiple type of sensing methods enable the automated detection and interpretation of potential illicit events occurring at wide zones. These approaches scale to large geospatial coverages and potential enable reasoning on detected events accordingly for establishing advanced situation awareness for safety and security management practitioners.

3.1 Physical Disturbance Event Detections at Fences of Critical Areas

Perimeter fences are widely used to protect Critical Areas such as water treatment plants, oil refinery, construction sites etc. Fence structures help to prevent only part of potential intrusions or postpone them. Therefore a high level of security is needed to monitor and investigate activities on and around fences. Accelerometers are relatively reliable tools which can be used for monitoring non rigid fences. While monitoring perimeter fences, two problems have to be addressed (1) Detect unusual events; and (2) classify these events to help with decision making and security related actions. Below, we present solutions to both of these problems using vibration sensors which are mounted on fences.

To efficiently detect unusual events along a perimeter fence, the developed event detection algorithm has to have the following properties: Fast, simple, and online; little or no interference of the user; data passed in small packets (1 or 2s of data); training stage enabled. Taking into consideration such requirements, an event detector based on Median Absolute Deviation (MAD) of signal and confidence interval method was adopted [5]. For each sensor directional axis, $y(N)$ will be a packet of data with specified window size N pass. The Median Absolute Deviation is a robust measure of data variability that can be calculated as follows:

$$MAD(y(N)) = \text{median}|y(N) - \text{median}(y(N))|$$

Then the median of this packet of data should be inside a confidence interval with a selected value range. The lower and upper bounds of the confidence interval are calculated as:

$$D(N)_{low} = \text{median}(y(N)) - \gamma * \sigma_y, D(N)_{up} = \text{median}(y(N)) + \gamma * \sigma_y$$

Where, the standard deviation of the signal in the given window is estimated as

$$\sigma_y = \frac{MAD(y(N))}{0.6745 * \sqrt{2}}$$

γ was selected to be equal to 4 in order to guarantee more than 99.7% of confidence in the samples to be within these bounds.

Then an event for a given axis is detected if $D(N)_{up} - D(N)_{low} > \text{threshold}$. The latter is estimated using packets of data when no-activities take place. The threshold indicates

an allowed deviation from the confidence interval when the packet of data will be considered while associated to no-activities. The quality of the MAD event detector is assessed using precision and recall measures. It was shown in the literature that such events as rattle, kick, climb or lean can be successfully classified for detecting security fence breaching under certain conditions [6]. In this paper, we will classify kick (K), shake (S) and no-activity (NA) events for each packet of data using a Bagging algorithm (Bag of decision trees) [7]. Cascade classification is also suggested in this paper. At the first stage, a packet of data is classified as Activity (A), No-activity (NA), Start (St) and End (E). If Activity (A) was classified, then this packet is classified as either kick (K) or shake (S). If the classifier returns Start event, then it is classified as a transition from no-activity to kick (NAK) or shake (NAS). The End event is classified as transition from kick or shake to no-activity (KNA, SNA). The initial investigation showed that misclassifications of K and S occur quite often during transition periods due to damping effects of the vibration signal. The quality of classification is assessed using Correct Classification Rate (CCR) for each state detection.

Experiments were performed using fence structure, as shown in Fig. 4. Each fence section 2 m high and 3 m wide. S1 and S2 indicate the locations of the vibration sensors.



Fig. 4. Schematic presentation of non-rigid fence and vibration sensors

Six tests were performed with 2 persons who kicked and shook various sections of the fence at different times. Overall, 30 kicks and 31 shakes were experimented and recorded. The sensors were left on the fence for 15 min to record no-activities which used to calculate a threshold for event detection (it was 0 in this case for both sensors). The sampling rate of sensors were at 100 Hz, and packets of 200 samples (2s of data) was passed to the MAD event detector and classifier. The start and end of events were manually labelled with some bias for the end of event due to the damping effect of the signal. For the training stage of the detector and classifier, 70% data was used for training (21 kicks and 22 shakes) while the rest of the data was used for testing (9 kicks and 9 shakes). In this paper, we consider High Level Event detection, i.e. an event is detected if the alarm was raised at least for one packet of data when the event takes place. Such event is marked as TP (True Positive). If the MAD event detector sets the alarm when no events took place, such event was marked as FP (False Positive). False Negative (FN) is counted when a whole event was missed. Table 1, below, shows recall and precision for individual tests and overall. Although all events were detected, some FPs occur usually in the end of an event. Examples of MAD event detector performances are shown in Fig. 5.

Table 1. Performance evaluation of MAD event detector

| Test | Test1 | Test2 | Test3 | Test4 | Test5 | Test6 |
|-----------|-------|-------|-------|-------|-------|-------|
| Recall | 100% | 100% | 100% | 100% | 100% | 100% |
| Precision | 93.3% | 72.7% | 100% | 84.2% | 88.9% | 90% |




Fig. 5. Performance of the MAD event detector: the best and average examples

| | | Predicted | | | | | | |
|--------|-----|--------------|--------------|--------------|--------------|--------------|--------------|-----------|
| | | K | KNA | NA | NAK | S | SNA | NAS |
| Actual | K | 87.01 | 0 | 0 | 0 | 11.69 | 1.3 | 0 |
| | KNA | 17.07 | 65.85 | 0 | 0 | 2.44 | 14.63 | 0 |
| | NA | 0 | 0.6 | 98.68 | 0 | 0 | 0.48 | 0.24 |
| | NAK | 9.09 | 0 | 0 | 63.64 | 0 | 9.09 | 18.18 |
| | S | 9.4 | 0.85 | 0 | 0 | 88.03 | 1.71 | 0 |
| | SNA | 6.52 | 13.04 | 6.5 | 0 | 10.89 | 63.04 | 0 |
| | NAS | 0 | 0 | 16.67 | 8.33 | 0 | 0 | 75 |

Fig. 6. Confusion Matrix for event classifications

A Cascade classifier which was suggested earlier achieved correct classification rate (CCR) at 93.75% overall. The Confusion Matrix for all detected events is given in Fig. 6.

NA is almost always classified correctly, the mis-classification usually occur during transition from K/S to NA which is expected due the nature of the vibration signals. NA is mis-classified as transition NAS sometimes when a packet of data contains more than 60% of NA samples. CCRs for both K and S exceed rates of 85%.

The results above showed that MAD event detector can be used to detect event reliably in Critical areas fences using vibration sensors, while the Cascade classifier can identify the nature of events taking place for further decision-support.

3.2 Unusual Behaviour Detection at a Toll Motorway

Automatic detection of incidents and unusual traffic events in motorways from visible spectrum videos is a challenging problem. These incidents range from: traffic collisions between vehicles and between vehicles and road structures, cars driving in the opposite direction or reversing, pedestrians and animals crossing the motorway, and more. Two

main approaches are considered for this problem. First approach is through learning the environment and thus learning which motions and behaviours are usual in this environment [8–11]. The second possibility is the direct approach of detecting the specific objects and further detecting their motion and appearance and finally classifying these in terms of behaviours and events. Works on object classification and tracking [12], human detection and tracking [13–15] and human behaviour recognition [16] falls into this category.

Given the diversity of incident types, we have initially opted for the more generic approach of learning the usual behaviour/motions of the scene. Please note that this choice is often a trade-off between the flexibility and accuracy of the detector. Furthermore, the number of training examples are often limited and this would hinder the design of specialized detectors. A hybrid approach has since been developed to handle a specific case of stationary cars in a tunnel and a more generic detector for the non-roofed areas.

Detecting Unusual Behaviour via Learning the Usual Flow

The approach here is similar to the method introduced by Adam et al. [8] where a grid of local monitors learn the low-level local statistics of physical motions. A monitor will produce a local alert if the observed motion does not conform with the usual patterns of motion in that neighbourhood. These alerts are then fused across spatio-temporal windows to make the decision regarding the existence of an unusual event. The hypothesis states that incidents are events that disrupt the usual traffic flow in a motorway; and therefore can be detected as samples that do not fit the modelled usual flow. Figure 7 summarizes this method. Two examples of detected unusual behaviour are also shown, where the area with unusualness has been highlighted automatically. These two examples show a car driving in opposite direction and a dog crossing the motorway.




Fig. 7. Unusualness detection on traffic flow in motorways

Detecting Stationary Vehicles in a Tunnel

The specific problem considered here concerns the detection of a stationary car and a pedestrian on the pavement in a tunnel while the traffic is in a one-way flow. The placement of the camera is such that the images of the vehicles are captured from a side/frontal view. It was found that the accuracy of detection using the above method is low due to some inherent difficulties of the set. These difficulties include: (i) The specific pose and car headlights, which give rise to a significant amount of erroneous motion detections using optical flow; (ii) Motion of cars in the left lane of the road are near to parallel to the camera's principal axis, due to direction of travel and the placement of the camera.

As a result, the optical flow values of the stationary car do not produce the required signal to noise ratio for detection. A combination of background subtraction methods and a blob tracker is used to detect the stationary car and the person on the pavement. In this, the temporal variance-based method introduced by Joo and Zheng [17] and the median background subtraction are combined to obtain the robustness of temporal variance and capability of the median model to detect stationary objects. Further morphological transforms are used to clean the foreground, in order to assist the detection of distinctive blobs in the foreground. The detected foreground blobs are compared between two consecutive frames based on the size and motion of the blobs using a Kalman filter. The outcome of tracking is shown in Fig. 8.



Fig. 8. Blob tracking based on foreground detection

4 On Site Integration Pilot (OIP) and Future Development

In December 2016, a prototype demonstrator was implemented on-site at an ATTIKES toll motorway, Athens, Greece. This prototype consisted of the major elements with reliable messaging using RabbitMQ and algorithm processing codes which were implemented within linux containers (Docker). The processing was then passed to a reasoning engine which performed event stream processing and high level fusion to make final assertions about the sensor processing and detected events in real-time. Several days of testing were performed along with a final live exercise evaluated by end-users. The use of containers made fixing and redeploying the detection algorithms very efficient. The aim was for the Motorway Traffic Management Centre at Attikes to be enabled with automated detection of unusual events in their motorways sections. With tens of surveillance cameras in place at the Centre, staff cannot efficiently detect all unusual events with rapid response. The SDAIM consequently provides real-time alerts to events for staff which will reduce their response time lags by 50% according to the discussion which we conducted with Attikes staff. Additionally, and due to the flexible nature of the

SDAIM architecture, we will deploy further performing algorithms for unusual event detection. The final version of the SDAIM system will be tested and validated at the Traffic Management Centre for Attikes Motorway in May 2018.

Acknowledgement. The authors would like to thank partners TEKNIKER and Attikes for the acquisition of experimental data using accelerometers and CCTV respectively. The ZoneSEC research project is partly funded by the European Union under contract Number: EC_FP7 607292.

References

1. ZONeSEC: Towards an EU framework for the security of wide-zones (2014–2018). <http://www.zonsec.eu/>
2. Lambert, D.: A blueprint for higher level fusion systems. *Inf. Fusion* **10**(1), 6–24 (2009)
3. Sabeur, Z.: Structured multi-level data fusion and modelling of heterogeneous environmental data for future internet applications. In: *Geophysical Research Abstracts*. EGU General Assembly 2013, vol. 15 (2013)
4. Zlatev, Z., Veres, G., Sabeur, Z.: Agile data fusion and knowledge base architecture for critical decision support. *Int. J. Decision Support Syst. Technol. (IJDSST)* **5**(2), 1–20 (2013)
5. Barat, V., Grishin, D., Rostovtsev, M.: Detection of AE signals against background friction. *J. Acoust. Emission* **29**, 133–141 (2011)
6. Yousefi, A., Dibazar, A., Berger, T.: Application of non-homogeneous HMM on detecting security fence breaching. In: *Proceedings of the ICASSP* (2010)
7. Meinshausen, N.: Quantile regression forests. *J. Mach. Learn. Res.* **7**, 983–999 (2006)
8. Adam, A., Rivlin, E., Shimshoni, I., Reinitz, D.: Robust real-time unusual event detection using multiple fixed-location monitors. *PAMI* **30**(3), 555–560 (2008)
9. Breitenstein, M.D., Grabner, H., Van Gool, L.: Hunting nessie-real-time abnormality detection from webcams. In: *IEEE 12th International Conference on Computer Vision (ICCV) Workshops* (2009)
10. Saligrama, V., Chen, Z.: Video anomaly detection based on local statistical aggregates. In: *Computer Vision and Pattern Recognition (CVPR)* (2012)
11. Yun, K., Kim, J., Kim, S.W., Jeong, H., Choi, J.Y.: Learning with adaptive rate for online detection of unusual appearance. In: *Bebis, G., Boyle, R., Parvin, B., Koracin, D., McMahan, R., Jerald, J., Zhang, H., Drucker, Steven M., Kambhampettu, C., El Choubassi, M., Deng, Z., Carlson, M. (eds.) ISVC 2014. LNCS, vol. 8887, pp. 698–707. Springer, Cham* (2014). https://doi.org/10.1007/978-3-319-14249-4_67
12. Shah, M., Javed, O., Shafique, K.: Automated visual surveillance in realistic scenarios. *IEEE MultiMed.* **14**(1), 30–39 (2007)
13. Dalal, N., Triggs, B.: Histograms of oriented gradients for human detection. In: *IEEE Computer Vision and Pattern Recognition* (2005)
14. Bouchrika, I., Carter, J.N., Nixon, M.S., Morzinger, R., Thallinger, G.: Using gait features for improving walking people detection. In: *International Conference on Pattern Recognition* (2010)
15. Niebles, J.C., Han, B., Fei-Fei, L.: Efficient extraction of human motion volumes by tracking. In: *IEEE Computer Vision and Pattern Recognition* (2010)
16. Chaquet, J.M., Carmona, E.J., Fernández-Caballero, A.: A survey of video datasets for human action and activity recognition. *Comput. Vis. Image Underst.* **117**(6), 633–659 (2013)
17. Joo, S., Zheng, Q.: A temporal variance-based moving target detector. In: *IEEE International Workshop on Performance Evaluation of Tracking and Surveillance (PETS)* (2005)