MOSAIC
Multimodal Analytics for the Protection of Critical Assets

Atta Badii[1], Marco Tiemann[1], Richard Adderley[2], Patrick Seidler[2], Rubén Heras Evangelio[3], Tobias Senst[3],Thomas Sikora[3], Luca Panattoni[4], Matteo Raffaelli[4], Matthew Cappel-Porter[5], Zsolt L. Husz[5],Thomas Hecker[6] and Ines Peters[6]

[1]*University of Reading, Whiteknights, Reading, RG6 6AH, U.K.*

[2]*AE Solutions (BI) Ltd., Badsey, U.K.*

[3]*Technische Universität Berlin, Einsteinufer 17, 10587 Berlin, Germany*

[4]*Synthema srl, Pisa, Italy*

[5]*BAE Systems, Advanced Technology Centre, Bristol, BS34 7QW, U.K.*

[6]*DResearch Digital Media Systems GmbH, Otto-Schmirgal-Str. 3, 10319 Berlin, Germany*

Keywords: Multimodal Analytics, Decision Support, Video Analytics, Text Mining, Social Network Analysis, Security, Critical Assets, Critical Infrastructure, UI-REF.

Abstract: This paper presents an overview of the MOSAIC architecture and the validated Demonstrator resulting from an EU-co-funded research project concerned with the development of an advanced system for the use and integration of multimodal analytics for the protection of critical assets. The paper motivates the MOSAIC vision and describes the major components of the integrated solution; including the ontological framework, the data representation, text mining, data mining, video analytics, social network analysis and decision support. In the descriptions of these components, it is illustrated how MOSAIC can be used to improve the protection of critical assets without necessitating data gathering that goes beyond what is already currently being gathered by relevant security organisations such as police forces by improving data analytics techniques, integration of analysis outputs and decision support mechanisms.

1 INTRODUCTION

Gathering and analysing available data is the critical process that is necessary in order to protect critical assets as well as in many other tasks that are carried out by security organisations such as police forces. The EU-co-funded research project MOSAIC provides such organisations with the means to improve situational awareness when protecting critical assets, which may include diverse types of assets such as security-relevant buildings and large gatherings of persons that need to be protected. MOSAIC applies automated multimodal analytics techniques over data that is currently available within police forces, including database data such as conviction records, unstructured text data such as police notes and reports and video data from CCTV camera systems. The goal of MOSAIC is to reduce the need for labour-intensive and time-consuming manual data gathering and processing that is still common practice in police organisations around Europe. It is expected that this will enable, in particular, intelligence analysts and CCTV system operators, two key target user groups of MOSAIC, to focus on tasks that cannot or should not be automated.

This paper gives a high-level overview over the MOSAIC system as a whole and the constituent components of which it is composed. Sections 2 and 8 are concerned with introducing and summarising the contribution of MOSAIC as an overall system. Sections 3, 4 and 5 introduce text and data mining,

Badii A., Tiemann M., Adderley R., Seidler P., Heras Evangelio R., Senst T., Sikora T., Panattoni L., Raffaelli M., Cappel-Porter M., L. Husz Z., Hecker T. and Peters I..
MOSAIC - Multimodal Analytics for the Protection of Critical Assets.
DOI: 10.5220/0005126503110320
In *Proceedings of the 11th International Conference on Signal Processing and Multimedia Applications* (MUSESUAN-2014), pages 311-320
ISBN: 978-989-758-046-8

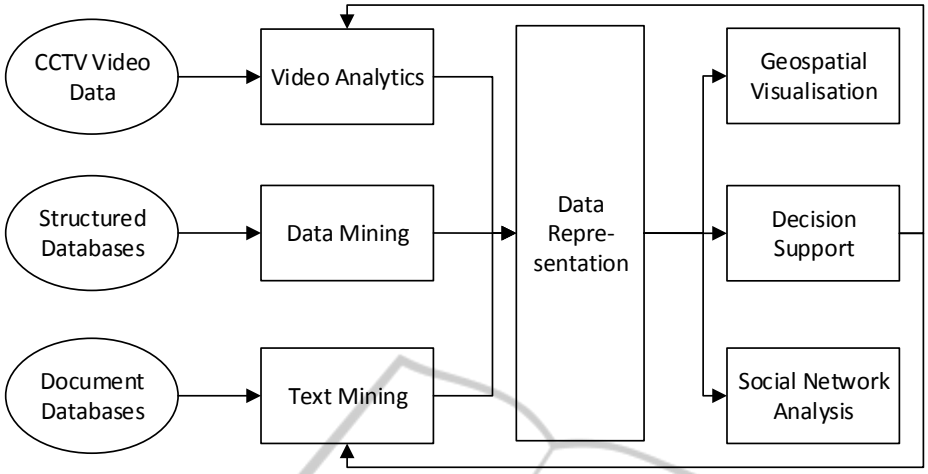

Figure 1: Overview of main MOSAIC system components.

video analytics and data representation components of the system that assist with the tasks of analysing available raw text and video data and making the analysis results available in a unified system. Sections 6 and 7 describe the criminal network analysis and decision support system functionalities that use the available and integrated data in order to assist analysts and CCTV system operators with analysis and visualisation functionalities.

2 MOSAIC VISION AND ARCHITECTURE

2.1 Vision

The vision for the MOSAIC system addresses several complementary goals.

First, it sets out to support end users that are tasked with evaluating a large amount of available and/or incoming data that are currently not accessible via a unified system infrastructure in carrying out their work more efficiently.

Second, it aims to enable organisations to automatically analyse incoming data in order to reduce end user cognitive load and in order to reduce the volume of network traffic required in particular for CCTV video surveillance involving high-definition video content transmission.

Third, it aims to support end users in carrying out data analyses without creating a fully automated process of decision making in order to retain human control, supervision and accountability while improving the effectiveness of operations supported by MOSAIC.

The vision described above is motivated by the current state of the market observed in organisations that are tasked with protecting critical assets – in particular policing organisations. Data available to such organisations is often fragmented, poorly indexed and analysed as well as not accessible in a way that facilitates efficient usage of individual data sources and in particularly not the timely discovery of relevant information across data from different data sources. MOSAIC aims to demonstrate how these issues can be overcome using data sources and infrastructure as they are available to these organisations today and taking account of realistic requirements concerning legal frameworks and conditions for data collection, retention, access and usage.

Existing and proposed solutions in the domain of concern for MOSAIC generally focus on a specific modality of input data such as video only (e.g. in Francisco et al., 2007; Bocchetti et al., 2009). These systems do not focus on integrated multimodal solutions, nor are open to data input in a simplified text format from heterogeneous sources independent (e.g. in Li, 2011). A system that provides somewhat similar functionalities as the MOSAIC system is the "Domain Awareness System" announced by Microsoft and the New York Police Department in late 2012 (Microsoft, 2012). Exact capabilities of the system as well as the deployment status of the system are however not publicly known at the time of writing.

2.2 Architecture

MOSAIC consists of components that address three main areas:

- Efficient processing and data analysis, using video analytics, text mining, data mining and distributed "smart edge" video processing;
- Unified data representation, storage and access, using Semantic Web technologies for unified representation and access;
- CCTV operator and intelligence analyst decision support, using data visualisation and geo-localisation, data analytics and production rule engines.

Components in each of the three areas are implemented as independent software systems. The system components are loosely coupled and communicate using SOAP and RESTful Web Services depending on suitability and security requirements. The architecture, design and validation of the MOSAIC system is supported by the UI-REF Framework which has provides the basis for methodologically-guided holistic requirements elicitation and prioritisation including the resolution of socio-ethical, legal and security requirements (Badii, 2008). This has underpinned the design of all use-cases as well as data representation, modelling and validation of system performance. Figure 1 depicts a summary overview of the individual components of the MOSAIC system.

MOSAIC uses text and data mining to integrate realistic example police information databases containing structured information (e.g. databases of criminal convictions) as well as ones containing unstructured information (e.g. full-text police officer notes describing observations made), and MOSAIC uses video analytics to identify trajectories and events in CCTV security video camera footage. A central data storage component represents data extracted using all of these methods using Semantic Web standards. A criminal social network analysis tool, an advanced geospatial visualisation system and data-driven decision support functionalities support system users and can be used to suggest specific actions to the system users where appropriate.

3 TEXT AND DATA MINING

3.1 Text Mining

The Text Mining (TM) component identifies relevant knowledge from sanitised Police reports and from sanitised free text fields in crime-related databases, by detecting entities and entity relationships. Named Entity Recognition (NER) and TM are applied through a pipeline of linguistic and semantic processors that share a common knowledge base with crime patterns, abbreviations, police terminology, acronyms and jargon. The shared knowledge base guarantees a uniform interpretation layer for the diverse information from different sources.

The automatic linguistic analysis of textual documents is based on morpho-syntactic, semantic, Semantic Role Labelling (SRL) and NER criteria. At the heart of the TM system is McCord's theory of Slot Grammar (McCord, 1980; McCord, 1990). The system analyses each sentence, cycling through all its possible constructions and trying to assign the context-appropriate meaning – the "right" sense – to each word. Each slot structure can be partially or fully instantiated and can be filled with representations from one or more statements to incrementally build the meaning of a statement. This includes most of the treatment of coordination, which uses a method of "factoring out" unfilled slots from elliptical coordinated phrases. The parser – a bottom-up chart parser – employs a parse evaluation scheme for pruning away unlikely analyses during parsing, as well as for ranking final analyses, which incrementally builds a syntactical tree. By including the semantic information directly in dependency grammar structures, the system relies on semantic information combined with semantic role relationships (*agent*, *object*, *where*, *when*, *how*, *cause*, etc.). The Word Sense Disambiguation (WSD) algorithm also considers possible super-subordinate-related concepts in order to find the appropriate senses in lemmas being analysed.

Appropriate heuristics have been implemented for MOSAIC in order to identify specific pre/suffixes, linguistic patterns and data formats for the English language so as to recognise key entities in text: dates, addresses, person names, locations, licence plate numbers, brands and manufacturer names, web entities, bank accounts and phone numbers. Once entities and semantic roles have been retrieved, the TM component then extracts entity relationships from the text. Two entities that are linked by a direct relationship such as agent-object/agent will have a very strong bond. Entities that have a relationship of proximity are also extracted. This type of approach allows police analysts to discover important relationships between entities, even though these are not linked by a syntactic dependency, for instance a person's name followed by a phone number in parentheses – proximity in the sentence – or a date and the author

at the beginning and the end of a document respectively – proximity in the document.

3.2 Data Mining

The purpose of the Data Mining (DM) component is to elevate typical tasks in the elicitation and preparation of data from disparate data sets and to implement algorithms for data analysis and the provision of results that are immediately and easily applicable in the intelligence analysis cycle.

On the data level, poor data quality is a constant issue in policing systems, resulting in suboptimal decisions. An entity resolution module has been developed in order to attend to those issues by determining whether records on persons refer to the same individual. This engine with predefined probability levels for matches on specific database field types is used to calculate a final match probability through a Bayes function. Optionally, fuzzy string matching is available through Metaphone (Philips, 1990) and Soundex (Odell, 1956) implementations. Preliminary results on the performance and accuracy of the algorithm in correctly matching entities show a minimum accuracy of 65% for 90% of test runs, whereas several combinations show an accuracy of up to 88% when compared to a Gold Standard test dataset. For comparison, compilation of the Gold Standard took the analyst 1½ days, involving handcrafting 1002 offender records into sets containing the same individual.

Data can be imported into the DM workbench which integrates DM algorithms for the access, preparation and analysis of data. As current intelligence models do not provide a structured approach to DM tasks, analysts can use data search and linking, exploration, modelling and visualisation capabilities through a process of interconnected nodes, following the formalisation of the Cross-Industry Standard Process for Data Mining (CRISP-DM; Shearer, 2000) model in conjunction with the intelligence cycle. The approach taken accommodates the various possible working environments and data requirements in which the final system could be applied. The resulting DM processes are reusable and can be re-run any time taking into account newly arrived data.

Together with police analysts, workbench functionalities have been tailored to the specific domain needs, resulting in three data mining processes:

- Offender mining and automatic assignment of priorities to offenders: tracking of prioritised criminal behaviour to enable law enforcement to allocate responsive actions in order to meet police priorities.
- Identification of crime series and mapping of known offenders to unsolved crimes: application of self-organising maps to link spatial, temporal, modus operandi and overlay of offender data onto clusters of similar crimes, thereby suggesting possible involvement in crimes.
- Identification of criminal roles: offence based roles, group offenders by their role and apply a *k*-means clustering algorithm to determine the most prominent groups for all offenders.

These data mining processes can be accessed and configured by intelligence analysts using the DM workbench.

4 VIDEO ANALYTICS

The increasing number of cameras installed in large-area CCTV networks leads to a load problem in the respective CCTV networks. Even with modern highly efficient video compression algorithms, a high number of cameras (typically >100 cameras in, for example, a middle-sized train station) sending their video data to a central server means inherently a high data load on the network. The MOSAIC project offers a change of paradigm in video processing methodologies: instead of streaming raw images to a powerful, central processing unit, each network node of a distributed network of smart cameras and video analysis units employs local processing and storage capabilities to translate the observed data into features and attributes. The major advantages of distributed processing are hence improved scalability and reliability of the network and better bandwidth utilisation. This network architecture has been implemented by adhering to the ONVIF specification (ONVIF, 2014). Following Senst et al. (2011), the video analytics sub-system consists of four classes of devices: (1) Network Video Transmitter (NVT) devices, which provide the video streams; (2) Network Video Analytic (NVA) devices, which analyse video, audio or metadata and provide the results in the form of metadata; (3) Network Video Display (NVD) devices, which represent media streams and the gathered information to human operators; (4) Network Video Storage (NVS) devices, which record the streamed video and its associated metadata. This network architecture is depicted in Figure 2.

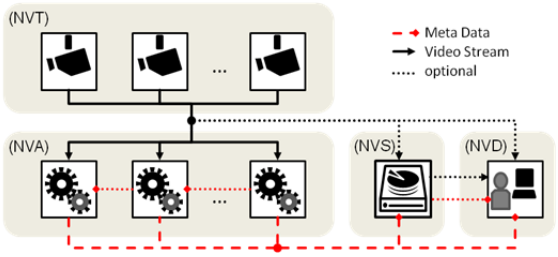

Figure 2: ONVIF-based network architecture.

Video analytics functionalities are provided by NVAs, which can accept video, audio or metadata information generated by other NVAs as input. This allows for a modularisation of the video analysis software. Three levels of video analysis are considered: (1) at the low level, features such as those provided by the detection of change (Evangelio et al., 2014) or the computation of optical flow between consecutive video frames (Senst et al., 2013), are computed; these features are used at the (2) mid-level in order to compute tracks of the objects of interest (Eiselein et al, 2012), i.e., persons (also across multiple cameras); (3) at the high-level, information gathered at the low and mid-level is used in order to extract semantic information such as crowd behaviour (Kuhn et al., 2012), multi-movement identification, human activities of interest (Acar et al., 2012), mugging detection, etc. The information extracted by each of the NVA entities is sent to the network as ONVIF messages, which can contain events and/or scene descriptions. Events are specific actions detected by ONVIF devices, usually having a semantic interpretation, to which a client can subscribe. Scene descriptions are XML-based abstractions of a scene in terms of predefined objects. The information provided by the whole set of NVAs is collected by the decision support and control sub-system for processing and display at different levels of abstraction. Furthermore, feedback and control information can be sent to the NVAs using Web Services.

Edge processing requires devices with suitable capabilities, in particular in terms of processing. This explains why until now only very simple algorithms have been ported to smart cameras - the performance of the smart cameras available on the market has so far not been sufficient for complex algorithms and procedures. These deficits are addressed in the MOSAIC project by developing hardware and firmware for multi-board intelligent IP cameras which feature sufficient computational power to facilitate on-board processing of state-of-the-art video analytics. The modular design of the MOSAIC smart camera comprises three boards which are connected by standard connectors: The System Board with internal and external interfaces and application specific auxiliaries, the Processor Board with the i.MX 6Quad processor, and the Sensor Board with the OV5640 system-on-a-chip (SOC) image sensor. This flexible design facilitates adaptation and exploitation.

Figure 3: MOSAIC Smart Camera. © DResearch, 2014.

One example of this is hardware-based detection of camera tampering events using the 3-axis accelerometer features of the MPU-9150 MotionTracking sensor on the system board.

Among video analytics functionalities considered in MOSAIC, the detection of left-behind objects is a clear example of a video analytics algorithm which can be brought to the edge of the CCTV network. The algorithm chosen for its implementation is presented by Evangelio et al in their paper (Evangelio et al., 2011). Modern foreground-background separation is implemented as well as multi-level modelling. Because long-duration scenes have to be modelled, high demands on local memory management arise. The adaption of these scenes to the complex lighting conditions results in considerable demands on parallel computing and floating point operations. Therefore, left-behind object detection algorithms have been chosen as

reference video analytics implementation which has been successfully ported to the MOSAIC smart camera.

5 DATA REPRESENTATION

The MOSAIC data representation component provides a central point of access for data available in the MOSAIC system. It encompasses a specific data format and data representation model and a software component that provides necessary functionalities to offer data store functionalities to other system components.

The MOSAIC data representation system represents output created by MOSAIC data analysis components using an ontological data representation model defined in the ontology representation format OWL-Lite (McGuinness and van Harmelen, 2004). The MOSAIC ontology represents relevant abstract concepts and their relations in a model. Data provided by the data extraction and analysis components is added to this abstract data model as instances that are realisations of concepts defined in the ontology model. All data is represented as 'triples' of entities and relations which form a directed acyclic graph. The embedding of the data instances into an ontological data model allows processing and reasoning that takes into consideration the entity properties and relations that are expressed in the ontology 'world model'. Figure 4 illustrates a basic benefit of using an ontology model for data representation.

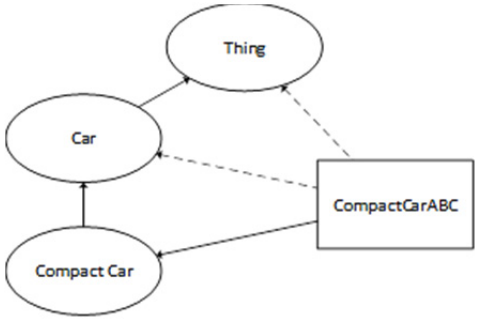

Figure 4: Example for implicit assertions through ontology modelling. The instance "CompactCarABC" is only explicitly declared to be of entity type "Compact Car", but since "Compact Car" is part of an ontological relation with "Car", "CompactCarABC" is also implicitly characterised as instance of "Car" (and of "Thing").

MOSAIC uses RDF to formally represent the data model internally as well as for communication with external components. RDF/XML is used for communication between components when exchanging data using the MOSAIC data model and format.

The MOSAIC data store provides users with functionalities to

- make the MOSAIC ontology and MOSAIC data accessible;
- create, read, update and delete data;
- query for data including queries involving implicit assertions as described above;
- subscribe to queries so that they receive notifications when new data lead to new results for the query they have subscribed to.

The MOSAIC data store is based on the Apache Fuseki system (Apache Fuseki, 2014) that itself uses the Apache Jena Semantic Web stack (Apache Jena, 2014) for data storage and representation. The MOSAIC data store extends Apache Fuseki with import/export functionalities, publish-subscribe functionalities and reasoning capabilities for ontology inferencing and making implicit assertions exploitable in queries. The MOSAIC data store uses the SPARQL query language for interacting with the stored RDF and OWL data. System functionalities are generally accessible via RESTful Web Services.

6 SOCIAL AND CRIMINAL NETWORK ANALYSIS

A criminal network generation, visualisation and analysis tool (Figure 5) enables the user to conduct social network analysis modified for the application onto criminal networks from data accumulated through the data and text mining components, facilitated by the semantically enabled data representation. The outlined techniques support analysts in hypothesis testing, which can be evaluated for accuracy within a domain specific, safe environment, and resulting operationalisation of its outcomes when data is turned into actionable intelligence. For generating co-offender networks, the method presented in Adderley et al., (2008) is applied. A multi-graph is established linking offenders associated with each other through one or more events (e.g. a crime) and this offender is also involved in events with other offenders. This process identifies many candidate networks which require prioritisation, i.e. edge, node and overall network weighting, to ensure that those who are causing the most harm are targeted first and that policing resources are allocated most effectively.

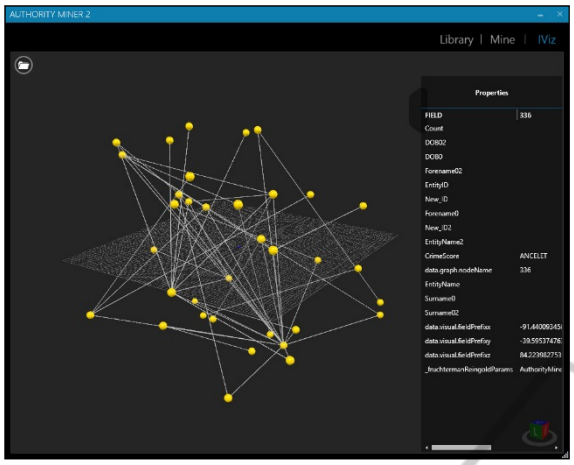

Figure 5: Sample network visualisation.

Weights are retrieved from social and human capital features. Social capital is estimated from the networks' centrality measures from the field of social network analysis, calculated, categorised and scored depending on the investigative context. Human capital comprises the individual contribution to the network. These include the individual's criminal role, travel distance to crime, and a domain based harm score. To retrieve criminal roles, high-level groups of crimes are created from aggregated crime types. If the frequency of a single crime type is above an adjustable threshold in the entire offender's crime set then it is labelled as a Significant Crime which is the role of that offender in the network. The harm score is based on its priority to the law enforcement agency to reflect current operational priorities. A weight is allocated to the age of a crime and is calculated by placing the crimes into date segments. Each segment is allocated a real number value between one and ten and used as a multiplier in conjunction with the crime score to assign a harm significance value to the crime as a prioritisation criterion. The distance that an offender travels is used to identify whether the person offends close to home (a known and familiar area) or commits crimes through a wider area (less known possible meaning a different type of person). Merging that distance with the main offences that the offenders commit and assessing the average and standard deviations for every crime committed by an offender we can assign an offender's significance given the specific context of the investigation.

All significant attributes for the definition of capital in the network can be combined into a definition of a criminal profile as presented in Table 1. Prioritised profiling is achieved by combining all scores into a final score that can now be allocated to each offender and thus a prioritised list of offenders is retrieved (see Table 2 for sample) which facilitates decision making in targeting the appropriate person(s) based on their effective final score(s).

Table 1: Criminal profile definition.

Profile attribute	Definition
Criminal Role	Criminal's crime type preference from an aggregate of her/his recorded crimes
Travel Distance	Average over the Euclidean distances between offender home address and crime location
Harm score	Crime score weighting: sum of scores of all crimes the offender has been involved in
Criminal activity	Indicated by the degree centrality, i.e. the number of links a vertex has
Information control	Indicated by the between-ness centrality, also known as gatekeeper function
Access	Indicated by the closeness centrality

To further evaluate network dynamics, analysts are provided with possibilities to investigate visually the evolution of a network. A sliding window filter combined with a sampling period (Kaza et al., 2007) is used to capture the network at significant states of its evolution, an evolutionary modelling approach superior to random sampling. Moreover, simulation techniques can be applied to investigate potential intervention strategies for the disruption of a criminal network and evaluate their effectiveness through both the structural and strategic damage to the network.

7 DECISION SUPPORT

Beyond the support for data analysts provided by the data store and the social and criminal network analysis component, the MOSAIC decision support components provide functionalities that can be used by intelligence analysts as well as CCTV operators and functionalities specifically targeting situational awareness in real-world environments which can be expected to be particularly beneficial to CCTV operators.

Table 2: Example for Top 3 Criminals by Force-prioritised profile scores.

ID	Role	Harm	Distance	Inform.	Access	Activity	Score
1	Burglary	600	Compact	Controller	Best	Active	30.49
2	Burglary	600	Compact	None	Best	None	27.49
3	Violence	360	Compact	None	Average	None	17.91

The relevant decision support functionalities can be grouped into the categories of Decision Support Engines and Information Visualisation. Both are described in the following subsections.

Two Decision Support Engines have been implemented as part of the MOSAIC system: a Template Matching Decision Support (TMDS) tool, developed by BAE Systems, and an Ontology-Based Production Rule (OBPR) system, developed by the University of Reading.

The TMDS tool allows low-level events to be grouped together into higher-level events to help predict when a crime is about to happen. The tool responds to real-time events generated by the Networked Video Analytics and matches patterns amongst the newly formed combination of these events and those already known. The TMDS focuses on analysing data events received from the Networked Video Analytics components, but can also integrate input from other data sources via the MOSAIC data store. The TMDS is configured with easy-to-write user-defined templates of occurrences of information and/or observed events. Templates can be constrained by metadata such as geographic area, time window, or other event properties (e.g. the detected vehicle registration number). For each scheduled rule, the TMDS automatically carries out user-specified actions, such as generation of new MOSAIC system events and automatic tasking of CCTV cameras. Higher-level events that are generated as an output of the matching process feed-back to the TMDS and the MOSAIC data store can be used as the input of another template that allows for further "meta-level" matching and processing.

The OBPR component generally carries out the same tasks as the TMDS tool, but works directly on the MOSAIC data store and offers different features and constraints. The OBPR component uses the mature JBoss Drools rule engine (JBoss Group, 2014) for production rule processing and the associated powerful JBoss Drools Rule Language for the formulation of decision support rules. In addition, the rule engine has been integrated with the data store so that it can directly use the MOSAIC ontology model in formulated rules. This enables the OBPR component to formulate complex rule types involving for instance the temporal order of events observed, or to evaluate rules that use the world model described in the MOSAIC ontology. The system can also support a wide range of actions to be taken when a rule 'fires', such as submitting control commands to an ONVIF CCTV network for approval by an operator, dispatching different types of notification messages or amending the MOSAIC data model with additional derived instances.

In addition to relational graph data visualisation, MOSAIC incorporates an advanced geospatial visualisation system. The VIKI (View It, Know It) visual situation awareness tool (Figure 6) provides an interactive mixed reality environment which displays live or recorded data, including live video feeds. These are projected into a synthetics 3D world representative of an observed area. The technology has been developed by BAE Systems, based on CAST Ltd. Viewers' Situational Awareness for Applied Risk and Reasoning (VSAR) toolkit.

In the MOSAIC system, VIKI acts as a CCTV Command and Control system in which the CCTV operator can "fly to the" areas of interest of the 3D visualisation or can perform a virtual patrol. CCTV video feeds are projected from virtual cameras so the operator assesses the feed seamlessly against the 3D model, other video streams and geo-located events. The operator gains visual feedback on camera positions, and can control pan-tilt-zoom CCTV cameras. The events from the video analytics (section 4) and decision support engines (section 7.2) are presented at their geo-location as soon as they are detected. The operator is able to view the low level events that led to this event, but also to view recorded video footage of these events.

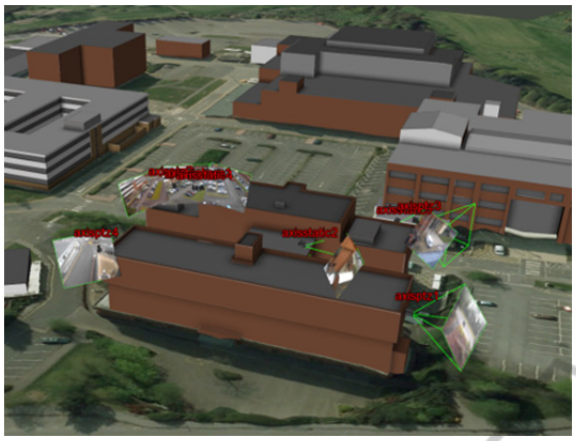

a.

b.

Figure 6: View It Know It (VIKI) system. a. Model of the MOSAIC site in VIKI b. List of geo-localized events and car watch list event marked by the blue sphere.

8 CONCLUSION

This contribution provides an overview of the system its components developed in the MOSAIC project in order to improve the protection of critical assets. MOSAIC in particular showcases the benefits of automated data processing in combination with a unified and integrated data representation. Advanced functionalities for data analysis and decision support show how these data can be employed with the aim of improving the speed and quality of information analyses and the situational awareness of system operators dealing with real-time data flows as well as with combinations of real-time and historical data.

MOSAIC shows how a complex end-to-end solution for the protection of critical assets can be created that can integrate the heterogeneous data that systems such as MOSAIC will be confronted with in real-world deployment situations. MOSAIC focuses on a limited set of key data sources, in particular selected police databases, written text reports and statements and CCTV video sources. The use and development of standard representation formats in MOSAIC facilitates the integration of similar as well as of different types of input data.

While an extension of the MOSAIC system with additional data sources is a potential area for future work, it is also important to investigate the impact of integrating available data, as carried out in MOSAIC, on the potential need for more, other, or less data sources. It may well be possible that improvements in the integration of available data may reduce the number of data sources needed in order to provide adequate protection. In turn this may minimise the amount of data collected in order to reduce the impact of data gathering on the privacy rights of persons who are potentially processed with a system such as MOSAIC.

ACKNOWLEDGEMENTS

The research leading to these results has received funding from the European Union's Seventh Framework Programme (FP7/2007-2013) under grant agreement no. 261776.

REFERENCES

Acar, E., Senst, T., Kuhn, A., Keller, I., Theisel, H., Albayrak, S., Sikora, T., 2012. Human Action Recognition using Lagrangian Descriptors. In: *IEEE MMSP*.

Adderley, R., Badii, A., Wu, C., 2008. "The Automatic Identification and Prioritisation of Criminal Networks from Police Crime Data," in *Intelligence and Security Informatics*, vol. 5376, D. Ortiz-Arroyo, H. Larsen, D. Zeng, D. Hicks, and G. Wagner, Eds. Springer Berlin / Heidelberg, pp. 5–14.

Apache Fuseki, 2014 http://jena.apache.org/documentation/ serving_data/.

Apache Jena, 2014. http://jena.apache.org/.

Badii, A., 2008. User-intimate requirements hierarchy resolution framework (UI-REF). in: *AmI-08*: *Second European Conference on Ambient Intelligence*, Nuremberg, Germany.

Boccetti, G., Fammini, F., Pragliola, C., Pappalardo, A., 2009. Dependable Integrated Surveillance Systems for the Physical Security of Metro Railways. Proceedings of the *Third ACM/IEEE International Conference on Distributed Smart Cameras, Como*, Italy, 30.8. - 2.9.2009.

Eiselein, V., Arp, D., Paetzold, M., Sikora, T., 2012. Real-time Multi-human Tracking using a Probability

Hypothesis Density Filter and Multiple Detectors. In: *IEEE AVSS*.

Evangelio, R.H., Paetzold, M., Keller, I., Sikora, T., 2014. Adaptively Splitted GMM with Feedback Improvement for the Task of Background Subtraction. *IEEE Transactions on Information Forensics & Security*, vol. 9, no. 5, May 2014, pp. 863-874.

Evangelio, R.H., Sikora, T. (2011). Complementary Background Models for the Detection of Static and Moving Objects in Crowded Environments in: Proceedings of the *IEEE International Conference AVSS.*

Francisco, G., Tillman, J., Hanna, K., Heubusch, J., Ayers, R., 2007. Integrated Homeland Security System with Passive Thermal Imaging and Advanced Video Analytics. Proceedings of the *SPIE Conference on Infrared Technology and Applications, Orlando, Florida*, 9.3.2007.

Kaza, S., Hu, D., Chen, H., 2007. Dynamic social network analysis of a dark network: Identifying significant facilitators. *Intelligence and Security Information, 2007 IEEE.* Pp. 40-46. IEEE.

Kuhn, A., Senst, T., Keller, I., Sikora, T., Theisel, H., 2012. A Lagrangian Framework for Video Analytics. In: *IEEE Workshop on Multimedia Signal Processing.*

Li, X., 2011. Design and Implement of Decision Support System for Police Emergency Response. Proceedings of the *3rd International Conference on Computer Research and Development, Shanghai, China*, 11-13 March 2011.

McCord, M.C., 1980. Slot Grammars. American Journal of Computer Linguistics. 6, 31–43.

McCord, M.C., 1990. Slot Grammar: A System for Simpler Construction of Practical Natural Language Grammars. Proceedings of the *International Symposium on Natural Language and Logic*. pp. 118–145. Springer-Verlag, London, UK.

McGuiness, D., van Harmelen, F. (2004): OWL Web Ontology Language Overview. W3C Recommendation 10 February 2004. http://www.w3.org/TR/owl-features/

Microsoft, 2012. Microsoft and MYPD Announce Partnership Providing Real-Time Counterterrorism Solution Globally, Press release.

Odell, M. K., 1956. The profit in records management. Systems, 20, 20.

Open Network Video Interface Forum, 2014. http://www.onvif.org.

Philips, L., 1990. Hanging on the Metaphone. Computer Language, 7, 12, 39-43.

Senst, T., Geistert, J., Keller, I., Sikora, T., 2013. Robust Local Optical Flow Estimation using Bilinear Equations for Sparse Motion Estimation. *In: 20th IEEE International Conference on Image Processing*.

Senst, T., Paetzold, M., Evangelio, R.H., Eiselein, V., Keller, I., Sikora: T., 2011. On building decentralized wide-area surveillance networks based on ONVIF. In: Proceedings of the *IEEE Conference AVSS.*

Shearer, C., 2000. The CRISP-DM Model: The New Blueprint for Data Mining. Journal of Data Warehousing, 5, 4, 13–22.