

9

Complex Project to Develop Real Tools for Identifying and Countering Terrorism: Real-time Early Detection and Alert System for Online Terrorist Content Based on Natural Language Processing, Social Network Analysis, Artificial Intelligence and Complex Event Processing

Monica Florea¹, Cristi Potlog¹, Peter Pollner², Daniel Abel³, Oscar Garcia⁴, Shmuel Bar⁵, Syed Naqvi⁶ and Waqar Asif⁷

¹SIVECO Romania SA, Romania

²MTA-ELTE Statistical and Biological Physics Research Group, Hungary

³Maven Seven Solutions Zrt., Hungary

⁴Information Catalyst, Spain

⁵IntuView, Israel

⁶Birmingham City University, United Kingdom

⁷City, University of London, United Kingdom

E-mail: Monica.Florea@siveco.ro; cristi.potlog@siveco.ro;

pollner@angel.elte.hu; daniel.abel@maven7.com;

oscar.garcia@informationcatalyst.com; sbar@intuview.com;

Syed.Naqvi@bcu.ac.uk; Waqar.Asif@city.ac.uk

In the last decades, the importance of social media has increased extremely with the creation of new communication channels and even changing the way people are communicating. These trends came along with the disadvantage of allowing a new scenario where messages containing valuable data about critical threats like terrorism and criminal activity are ignored, due to the

sheer inability to process – much less analyze – the vast amount of available data. Terrorism has a very real and direct impact on basic human rights of victims, such as the right to life, liberty and physical integrity, often with devastating consequences.

In this context, the RED-Alert project was designed to build a complete software toolkit to support LEAs in the fight against the use of social media by terrorist organizations for conducting online propaganda, fundraising, recruitment and mobilization of members, planning and coordination of actions, as well as data manipulation and misinformation. The project aims to cover a wide range of social media channels used by terrorist groups to disseminate their content which will be analysed by the RED-Alert solution to support LEAs to take coordinated action in real time but having as a primordial condition preserving the privacy of citizens.

9.1 Introduction

Radicalisation leading to violent extremism and terrorism is not a new phenomenon but the way it is now spreading is more and more alarming and extending to the EU as a whole. As a matter of urgency, the European and Member States' policies must evolve to match the scale of the challenge offering effective responses [1].

During recent years Europe is facing new challenges to design and build new tools and to take advantage of technological advancements to prevent terrorist attacks. The Europol report from 2017 shows that, in 2016, a total of 142 failed, foiled and completed attacks have been reported. In 2017, 16 attacks struck eight different Member States while more than 30 plots were foiled.¹

The RED-Alert project is aligned to SECURITY Work Programme 2016–2017 call objectives that targets improvement of investigation capabilities, solving crimes more rapidly, reducing societal distress, investigative costs and the impact on victims and their relatives and to prevent more terrorist endeavours.²

The RED-Alert project is a H2020 European research and development project that uses analytics techniques such as NLP, SMA, SNA and CEP to

¹https://ec.europa.eu/home-affairs/sites/homeaffairs/files/what-we-do/policies/european-agenda-security/20180613_final-report-radicalisation.pdf

²http://ec.europa.eu/research/participants/data/ref/h2020/wp/2016_2017/main/h2020-wp1617-security_en.pdf

tackle LEAs needs in terms of prevention and action regarding terrorist social media online activity.

The novelty the project brings is combining these technologies for the first time in an integrated solution that will be validated in the context of five LEAs.

The consortium was designed to gather together all required capabilities and expertise that sustain the development of RED-Alert solution:

Five Law Enforcement Agencies (LEAs): Protection and Guard Service from Republic of Moldova (SPPS), Guardia Civil from Spain (GUCI), Ministry Of Public Security – Israel National Police (MOPS-INP), Metropolitan Police Service from UK (SO15) and Protection and Guard Service from Romania (SPP);

Five Industrial innovation champions (of which four SMEs): SIVECO Romania SA (SIV), Intu-View Ltd (INT), Usatges Bcn 21 Sl (INSKT), Maven Seven Solution Technology (MAV), and Information Catalyt for Enterprise Ltd (ICE);

Four Academic & Research Organizations: Interdisciplinary Center Herzliya (ICT), Eotvos Lorand Tudomanyegyetem (ELTE), City University Of London (CITY) and Birmingham City University (BCU);

One Regulatory association: Malta Information Technology Law Association (MITLA).

The project duration is 36 months and started in June 2017.

9.2 Research Challenges Addressed

The main challenge in the domain of terrorism and radicalization research is that the underlying data sources and data usages are constantly and rapidly evolving, as terrorist groups are moving away from structured written blogs and forum posts and instead, are using social media to propagate URLs that redirect to repositories of propaganda videos. Thus, processes of detecting suspicious content can become quickly outdated, and it is becoming essential to automatically adapt the system to evolving media channels layouts and interfaces, as well as changing user behaviours.

To support the project's objectives, the following key performance indicators (KPIs) are to be reached until the end of the project: seven social media channels mined for content, 10 languages supported for analysis,

improved accuracy and usability of tools within the context of data privacy, as well as extended real-time and collaborative capabilities and support for further development.

To address these KPIs, the RED-Alert project mixes relevant software components from different partners. In the same time, the challenge and innovation are to combine technologies such as CEP, SNA and NLP to assess social features in communications used by terrorist organizations. This will imply harmonisation of theories, tools and techniques from cognitive science, communications, computational linguistics, discourse processing, language studies and social psychology. Moreover, in order that the system performance to be adapted for each component the project implements a meta-learning process that will assist SNA, CEP and NLP components defined processes.

Another major challenge that needs to be addressed by the project is to preserve the privacy of citizens that use online social networking platforms. Having in mind the rumours linked with social media data collection and new GDPR that applied from 25th of May 2018, became obvious that the Internet service providers struggle to balance the user privacy against the national security. The only way to move forward is to preserve the privacy when processing the data and in the same time to take advantage of the latest technological advancement when designing the security part of the system; hence, the malicious content and the corresponding personality can be tracked while the privacy of innocent citizens can be preserved. RED-Alert system will include privacy-preserving mechanisms allowing the capture, processing and storage of social media data in accordance with applicable European and national legislations.

RED-Alert will face the additional challenge of allowing collaboration between the different LEAs from different countries, with different privacy laws and trust levels by implement a privacy-preserving tool to mine the data.

There is a growing understanding that innovation, creativity and competitiveness must be approached from a “design-thinking” perspective – namely, a way of viewing the world and overcoming constraints that is at once holistic, interdisciplinary, integrative, innovative, and inspiring. Privacy, too, must be approached from the same design-thinking perspective. Privacy must be incorporated into software systems and technologies, by default, becoming integral to organizational priorities, project objectives, design processes, and planning operations [2].

9.3 Architecture Overview

The vision of RED-Alert project is to develop and validate a real-time system able to facilitate the timely identification of terrorism-related content by summarizing large volumes of data from social media and other online sources (such as blogs, forums).

The RED-Alert components as shown in Figure 9.1: NLP, SMA, SNA, CEP, Data Anonymization, Data Visualisation and ML) will be integrated in three separate layers, based on Lambda Architecture³ concepts defined by [3], designed to handle massive quantities of data by taking advantage of both batch and stream methods for real-time data processing, as follows:

- the “speed” layer, which includes data acquisition components for processing data streams in real time by means of data collection (social media capture, web crawling, LEA “raw” content), data filtering

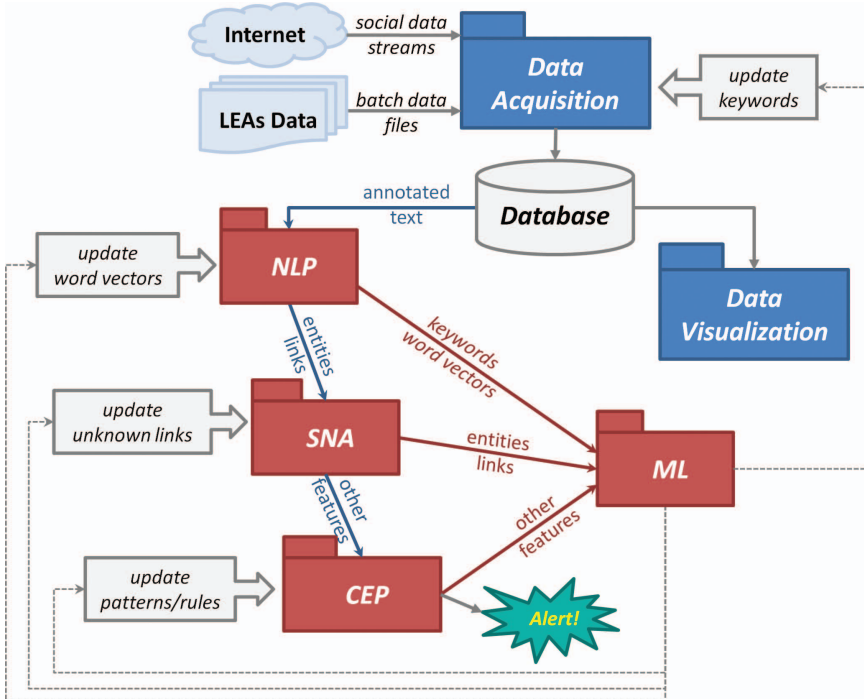


Figure 9.1 Dynamic learning capabilities of the systems to update keywords, vector spaces, rule patterns, algorithms and models.

³<http://lambda-architecture.net/>

(pulling text data from a message queue, normalizing and extracting the required meta-data), data enrichment (multimedia content analysis), and data privacy (anonymization of text and image data);

- the “batch” layer, which integrates the predictive models (based on CEP) that will be used by the pattern detection features within the analysis module. Due to the changing nature of the facts and behaviours, the set of stored models should be periodically re-trained with the new data arriving to the system. This is usually a resource-intensive task that cannot be performed in real-time and it should be scheduled as a frequent batch job;
- and the “service” layer, which integrates the visual analytics gateway that will be in charge of presenting the aggregated data, metrics and events configured by users, who can set up the rules or conditions for triggering alerts. This will be used directly by the rules engine to determine whether the conditions exist for a particular event type. The layer will also offer a Web Service API allowing third parties or LEAs in-house developers to build external components on top of the RED-Alert integrated solution.

Figure 9.2 shows the designed Architecture for RED-Alert. In this multi-layered architecture, application components are grouped into logical layers, namely:

- Front-end Layer – grouping components and functionalities that face the end-users of the system, with the role of getting and presenting data, displaying alerts, and allowing the users to configure the system and administrators to monitor it;
- Back-end Layer – grouping the core modules and data processing components that service the system;
- Integration Layer – grouping inward middleware services that inter-connect the components of the system, as well as outward facing APIs that facilitate connections with other systems;
- Data Storage Layer – grouping database management systems (both relational and non-relational) that handles the storage of data needed by the system.

This approach to architecture, described above, attempts to balance latency, throughput, and fault-tolerance by using batch processing to provide comprehensive and accurate views of historical operational batch data, while simultaneously using real-time stream processing to provide views of online data.

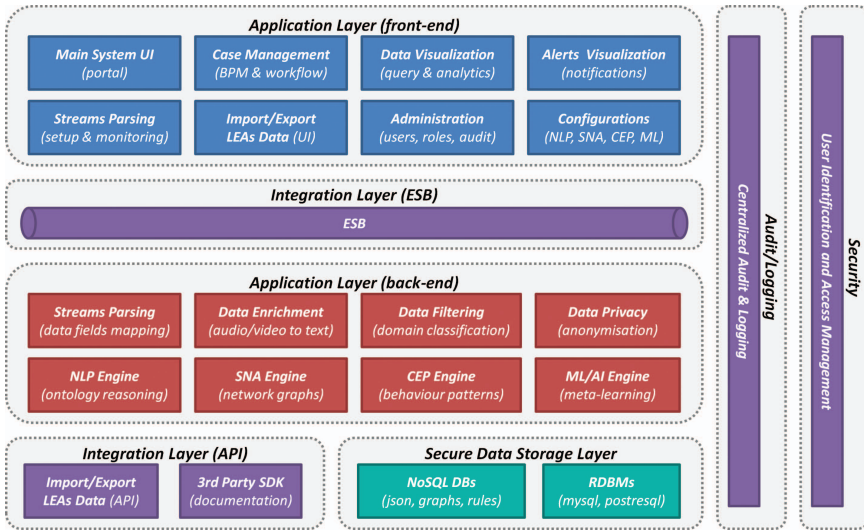


Figure 9.2 Layered application architecture.

The “speed” layer sacrifices throughput as it aims to minimize latency by providing real-time views into the most recent data. The “batch” layer pre-computes results using a distributed processing system that can handle very large quantities of data. Output from the “batch” and “speed” layers are stored in the “service” layer, which responds to ad-hoc queries by returning pre-computed views or building views from the processed data.

“Privacy by Design”, focuses on maximizing privacy and data protection by embedding safeguards across the design and development of software systems, services or processes by taking privacy and data protection considerations into account from the outset and throughout their whole lifecycle, rather than as a remedial afterthought. Such safeguards should be built into the core of the products, services or processes and treated as a default setting for not only technologies, but also into used operation systems, network infrastructures, work processes and management structures [4].

9.4 Results

9.4.1 Natural Language Processing Module (NLP)

Ever since the Tower of Babel, the human race has taken recourse to translation to bridge the gap between languages, cultures, societies and nations. Translation serves many purposes: it enables us to broaden the scope

of our cultural perspective, to see the world in a way that others – friends and foes – do, to retrieve ancient knowledge that, otherwise, would be lost to mankind and to communicate between people on a day to day basis.

However, in a global environment challenged with enormous amounts of information, a challenge has arisen that cannot be solved by translation. This is the need to identify affinities and dis-affinities between semantic units in different languages to normalize streams of information and mine the “meaning” within them regardless of their original language. When we look for information or wish to generate alerts – particularly in domains that are global – we do not want to be restricted to streams of information in one language; when we are interested in information – be it alerts on terrorism, fraud, cyber attacks or financial developments – we do not care if the origin is in English, French Arabic, Russian or Chinese. The need, therefore, is for technology that scans the entire gamut of information, identify the language and the language register of the texts, perform domain and topic categorization and match the information conveyed in different languages to create normalized data for assessment of the scope and nature of a problem.

The problem facing automated extraction of meaning from language is not restricted to translation between languages but within languages. That which we call a “language” is frequently a political definition and not one based on the linguistic reality. Some cases of a “language” are, actually a group of “dialects” that in other cases are defined as separate languages. The decision to call Swedish, Danish and Norwegian separate languages on one hand, and Moroccan, Libyan, Saudi Arabia and Egyptian all “Arabic” is political and not linguistic. Even within the same language register, words, quotations, idioms or historic references can be “polysemic”; they have different meanings according to the domain and the context of the surrounding text. A verse in the Quran may mean one thing to a moderate or mainstream Muslim and the exact opposite to a radical.

Methods to deal with this problem have generally been based on multi-lingual dictionaries that enable key words spotting (by input of a key word in one language, the search engine can add the nominal corresponding terms in other languages) or by automated translation of texts and application of the search criteria in the target language. The limitations of such methods are obvious: a word in one language has many “translations” and not all of them may even be remotely related to the meaning that the user is interested in.

In 1949 the cryptologist Warren Weaver wrote a memorandum on automated translation using computer technology. Weaver suggested the analogy, of individuals living in a series of tall closed towers, all erected over a

common foundation. When they try to communicate with one another, they shout back and forth, but cannot make the sound penetrate even the nearest towers. But, when an individual goes down his tower, he finds himself in a great open basement, common to all the towers. Here he establishes easy and useful communication with the persons who have also descended from their towers. Thus, he suggested "...to descend, from each language, down to the common base of human communication – the real but as yet undiscovered universal language – and then re-emerge by whatever particular route is convenient" [5]. In this description, Weaver touched – without calling it by name – on the approach that we are suggesting: semantic normalization of statements in different languages according to domain-specific ontologies.

This solution is based on emulation of the "intuitive" links that domain experts find between concatenations of lexical occurrences and appearances of a document and conclusions regarding the authorship, inner meaning and intent of the document. In essence, this approach looks at a document as a holistic entity and deduces from combinations of statements meanings, which may not be apparent from any one statement. These meanings constitute the "hermeneutics" of the text, which is manifest to the initiated (domain specialist or follower of the political stream that the document represents) but is a closed book to the outsider. The crux of this concept is to extract not only the prima facie identification of a word or string of words in a text, but to expand the identification to include implicit context-dependent and culture-dependent information or "hermeneutics" of the text. Thus, a word or quote in a text may "mean" something that even contradicts the ostensible definition of that text.

The meanings that are represented in one language by one word may be represented in other languages by completely different lexemes (words). "Idea Analysis" or "Meaning Mining" is the ability to extract from a text the hermeneutics (interpretation) that is not obvious to the non-initiated reader. We use of "Artificial Intuition" technology for this purpose. Artificial Intuition is based on algorithms that apply to input of unstructured texts the aggregated comprehension by seasoned subject matter experts regarding texts of the same domain used in training. Humans reach "intuitive" conclusions – even by perfunctory reading – regarding the authorship and intent of a given text, subconsciously inferring them from previous experience with similar texts or from extra-linguistic knowledge relevant to the text. After accumulating more information through other features (statements, spelling and references) in the text, they either strengthen their confidence in the initial

interpretation or change it. These intuitive conclusions are part of what the Nobel Laureate Prof. Daniel Kahneman called “fast thinking” – a judgment process that operates automatically and quickly, with little or no effort and no sense of voluntary control [6].

We have approached this problem through combining language-specific and language-register specific NLP with domain-specific ontologies. The technology extracts such implicit meaning from a text or the hermeneutics of the text. It employs the relationship between lexical instances in the text and ontology – graph of unique language-independent concepts and entities that defines the precise meaning and features of each element and maps the semantic relationship between them. As a result of these insights, the process of disambiguation of meaning in texts is based on a number of stages:

- Identification of the “register” of the language. The register of the language may represent a certain period of the language), dialects, social strata etc. In the global world today, however, it is not enough to identify languages; the world is replete with “hybrid languages” (e.g. “Spanglish” written and spoken by Hispanics in the US; “Frarabe” written and spoken by people of Lebanese and North African origin in France and Belgium) that are created when a person inserts secondary language into a primary (host) language, transliterates according to his own literacy, accent etc. It is necessary, therefore, to take the non-host language tokens, discover their original language, back transliterate them and then find the ontological representation of that word and insert it back into the semantic map of the document;
- Identification through statistical analysis (based on prior training of tagged documents) of the ontological instances in the text to determine the probability that the author represents a certain background and ideological leaning. Statistical categorization of a document as belonging to a certain domain, topic, or cultural or religious context can reduce the number of possible interpretations of a given lexical occurrence, hence reducing ambiguity;
- Disambiguation using the immediate neighbourhood of the lexical instances. Such neighbourhood consists of the lexical tokens directly preceding or following the lexical instance. After reading a number of texts of a given genre, the algorithm infers that X percent accord to statement A, the meaning B. When statement C is encountered in a text that is categorized as belonging to the same genre, the algorithm derives

from this a high level of confidence that C also means B. This confidence can be enhanced by additional information in the text;

- Statistical categorization of a document as belonging to a certain domain, topic, or cultural or religious context to reduce ambiguity;
- Chunking and Part of Speech Analysis of the text to use the relationship between different words (not necessarily arbitrarily choosing a certain level of N-grams) to provide additional disambiguating information;
- Based on the identification of the domain of the text, the lexical units (words, phrases etc.) are linked to ontological instances with a unique meaning (as opposed to words which may have different meanings in different contexts) that can be “ideas”, “actions”, “persons” “groups” etc. An idea may be composed of statements in different parts of the document, which come together to signify an ontological instance of that idea⁴;
- The ontological digest of the document then is matched with pre-processed statistical models to perform categorization.

This approach, therefore, is not merely “data mining” but “meaning mining”. The purpose is to extract meaning from the text and to create a normalized data set that allows us to compare the “meaning” extracted from a text in one language with that, which is extracted from another language.

This methodology applies also to entity extraction. Here, the answer to Juliette’s queclarative, “what’s in a name” is – quite a lot a not – as Juliette suggested almost nothing. A name can tell us gender, ethnicity, religion, social status, family relationships and even age or generation. To extract the information, however, we must first be able to resolve entities that do not look alike but may be the same entity (e.g. names of entities written in different scripts English, Arabic, Devanagari, Cyrillic) and to disambiguate entities that look the same but may not be (different transliterations of the same name in a non-Latin source language or culturally acceptable permutations of the same name).

⁴Ontology is a graph of unique language-independent concepts and entities built by experienced subject matter experts that defines the precise meaning and features of each element in the graph and maps the semantic relationship between them. Hence, the features that are encountered in the surroundings of a lexical instance are factored in the system’s decision to what unambiguous meaning (ontological instance) to refer the lexical instance. “Ontology”, Tom Gruber, *Encyclopedia of Database Systems*, Ling Liu and M. Tamer Özsu (Eds.), Springer-Verlag, 2009.

9.4.2 Complex Event Processing Module (CEP)

The key challenges so far with the complex even processing has been the need to make it both functional and generic at the same time. As downstream consumers, the component is dependent on receiving data from the other, upstream components, like the NLP and SMA data. The challenge here was to produce something that could consume unknown data as well as make assumptions and best guesses as to the nature, structure, quantity and quality of the data. In addition to working in the dark with its source data, the CEP engine also had to the challenge of not having any intelligence data to work with either. Clearly, on a project of this nature LEAs must guard and protect their intelligence for a plethora of operational reasons, however, regardless of this the CEP engine must still be delivered and demonstrate a working capability, so again it had to make a few leaps of faith which in the end should remain relevant when integration and pilot tests them out. Hence, the CEP engine remains probably a bit simplistic because of its generic nature, but by the same token generic is inherently extensible – so as both upstream data and real-world intelligence are fed into it, the engine will be able to adapt.

The CEP component aims to identify, via pattern matching algorithms, the dynamics, interactions, feedback loops, causal connections and trends associated with the data content it receives as input from the other RED-Alert components. Specifically, it is a secondary, downstream consumer of pre-processed data from the NLP, SNA and AI components and will generate output alerts; the component will also allow the configuration of data sources to allow the ingestion of external data out with the primary sources. The alerts themselves will be output to log files which will be monitored by a file reader component to display alerts, as well as monitor the CEP engine as a whole, and to integrate with the external API's of the LEAs.

As a development timeline, it comprises template architecture of many different CEP nuances which are set/selected/derived via a web tool to produce myriad applications. A data ingestion component will, either acquire processed data from the configured input component via configurable connection components, or the connection components will feed Kafka topics which will serve as the actual source for all CEP input. Apache Kafka⁵ is a distributed streaming platform generally used for two broad classes of applications: 1) for building real-time streaming data pipelines that reliably

⁵<https://kafka.apache.org/>

get data between systems or applications, and 2) for building real-time streaming applications that transform or react to the streams of data. And it is in this 2nd type of applications where the CEP from RED-Alert is conceived. Current expectations assume that multiple CEP applications will be running in parallel – each either working on different parts of the input data, or on different patterns within the data, or different configurations of the same CEP Application but utilizing an alternate configuration (e.g. on data consumed per month versus per week) or providing staged, partial result sets that will subsequently be consumed by an additional, downstream, CEP application that will act on the staged data.

Other tools and technologies covering similar RED-Alert needs and functionalities were analysed but dismissed as there was a need for extra developments further than the actual accomplished or the expertise of the development team was more limited. These other technologies were Apache Flink which is incorporated into the main CEP RED-Alert component, Spark or Red Hat Drools.

Primarily from a performance perspective, we expect Kafka to deal with this sort of load far better than Mongo, hence we expect any data sourced from Mongo to be moved into Kafka, and hence a data loading component will perform this task. Note also that as part of creating staged, pre-processed data for downstream consumption by other CEP applications, the CEP applications themselves will create and populate MongoDB/Kafka topics as well. Also, it is likely that Kafka will serve as the primary source for the engines and that these topics are populated from MongoDB, in real time. This event data is then converted to a data type associated with the CEP software via a generic parsing component to produce objects with a common structure representative of the source data (i.e. NLP, SNA and ML).

The block diagram shown in Figure 9.3 outlines the workflow, interactions, input/output and decision-making processes on the CEP engine itself. As the diagram clearly shows, the engine itself works on structured, well defined JSON, where well defined includes all field names, their data types as well as an indication of their original source – Note, in this case, source indicates where the data analytics (i.e. NLP, SNA and SMA processing) that generated particular aspects of the JSON originated, as opposed to the source of the input, i.e. the raw data.




Figure 9.3 Complex event processing module – Logical component diagram.

9.4.3 Semantic Multimedia Analysis Tool (SMA)

Multimedia is extensively used in social networks nowadays and is gaining popularity among the users with the increasing growth in the network capacity, connectivity, and speed. Moreover, affordable prices of data plans, especially mobile data packages, have considerably increased the use of multimedia by different users. This includes terrorists who use social media platforms to promote their ideology and intimidate their adversaries. It is therefore very important to develop automated solutions to semantically analyse given multimedia contents. The SMA Tool is designed to ensure security and policing of online contents by detecting terrorist material.

The SMA Tool extracts meaningful information from multimedia contents taken from social media. The five main features of the tool are:

- Segmentation of audio streams, identifying sections of speech;
- Transcription of the segmented speech sections using an ASR engine;
- Detection of sound events within audio streams, such as gunfire, explosions, crowd noise etc;
- Extraction and identification of objects, such as logos, flags, weapons, faces, etc., within image and video scene elements;
- Extraction and transcription of text elements in image and video elements.

Moreover, the SMA Tool retrieves multimedia data, converts it to a uniform format and delivers the analysis results. The extraction of semantic information is the third of four stages the tool will perform. All four stages are as follows:

- Input: Retrieval of multimedia files from disk or URL;
- Stream Separation: Extraction of audio/video streams in multimedia files;
- Feature Analysis: Semantic analysis of audio/image content;
- Output: Compilation of results in a uniform JSON format.

The results of this tool are sent to the other key components of the project such as NLP, SNA and CEP.

9.4.3.1 Speech recognition

This component is used for audio segmentation, language detection and speech transcription. The RED-Alert project is required to support 10 languages, and be able to run offline, without having to send data to a 3rd party web API. We have consulted our LEA partners to prepare a list of 10 languages which must be supported by the speech/written text transcription elements of the SMA Tool. These languages are: Arabic, English, French, German, Hebrew, Romanian, Russian, Spanish, Turkish, and Ukrainian.

9.4.3.2 Face detection

The SMA tool uses a Haar-like feature based cascade classifier [7] to detect both frontal facing and profile faces in images. Haar-like features are calculated by finding the difference in average pixel intensity between two or more adjacent rectangular regions of an image. In the SMA tool, Haar cascades are used as a supplementary feature to implement simple face detection. More advanced techniques are implemented in the object detection element, which can also be used to detect people/faces.

9.4.3.3 Object detection

State of the art methods for detection of objects within images use large neural networks consisting of multiple sub-networks (region proposal network, classification network etc.). The SMA tool's object detection utility uses the Faster R-CNN structure [8]. Faster R-CNN is constructed primarily of two separate networks: a Region RPN which produces suggestions of regions of an image which might contain objects, and a typical CNN which generates a feature map and classifies the objects in the proposed regions.

9.4.3.4 Audio event detection

Audio event detection is implemented in the SMA Tool by using a recurrent CNN [9]. The convolutional element classifies the short term temporal/spectral features of the audio, while the recurrent element detects longer term temporal changes in the signal. The SMA Tool applies feature extraction prior to processing by the network. This provides a more detailed representation of the audio signal to the network, meaning the first few layers can extract more meaningful information. Peak picking algorithms [10] are applied to remove any noise and only annotate the onset of any detected audio events.

9.4.4 Social Network Analysis Module (SNA)

In the last decades, human communication has gone through a crucial transition. Thanks to the Internet, which connects all individuals around the globe, everybody can contact each other without any time delay and without geographical restrictions. Social interactions became cheap and worldwide, the only restriction remained at the human side: all of us are able to process information at a finite rate and can engage trustful relations only with a few tens or hundreds of others. Therefore, describing and modelling of the new type of human interactions called for a description which is free of space limitations: these represent the tools of Network Science.

SNA module, aims to provide methods and software solutions for handling relational data. It focuses on three aspects of networked analysis as described in the following subsections.

1) Network dynamics and temporal network structure models.

The tool describes the evolution of networks and edges/nodes in time, by calculating quantitative features derived from models on evolving networks, and evolution of communities.

Real systems are usually not static, instead they evolve in time [11]. This can manifest in the emergence of new parts, the disappearance of existing parts, and also the relations among constituents can be rearranged over time. Temporal networks with changing topology over time result typically changing community structures. Since community finding methods determine the structures only at different time steps, the structures from consecutive steps must be matched. When communities simply shrink or increase in size, then the matching is straightforward: matching of communities is determined uniquely by intersecting nodes between the two communities of different time steps. However, individuals can also change their community membership over time.

The SNA module implements a special community finder algorithm to solve this challenge. The solution is based on the property of the applied algorithm, which ensures, that adding new nodes and edges to a network does not change the membership status of a node or an edge. The only possible change is, that distinct communities fuse. This property allows an algorithm to match consecutive groups by introducing an intermediate time step, where the two snapshots are merged into a common network. Because the intermediate snapshot can contain only additional nodes and edges, the communities of the intermediate network can be matched to the prior and to the subsequent communities by the rule of matching intersections.

2) *Link prediction solution*

The SNA tool of the RED-Alert solution adopts network theoretic similarity and distance measures for counter terrorism purposes. Based on the special targeted measures, missing links and nodes are predicted by the module. Furthermore, some features e.g. weights, labels, directionality of the links are updated as well.

The implementation relies on two theoretic pillars:

- prediction based on topological measures;
- prediction based on attribute information.

Topological measures use only information from connectivity patterns, in contrast attribute measures predict missing/hidden relations from common attribute statistics. Upon request of the analyst on the user interface of the integrated RED-Alert solution, the SNA module can apply hybrid predictions as well, where both networked measures and attribute data are combined (Soundarajan & Hopcroft, 2012).

It must be noted though, that all theoretical speculations are useless without reliable data sources. The scientific background behind this tool ensures only the mathematical rigour with the calculations, but the final conclusions must be always thoroughly reviewed by human experts. All mathematical models work with assumptions that can be only partially valid in real scenarios.

3) *Hierarchy reconstructing methods*

Terrorism has its own frame and structure. As all organizations that consist of many individuals and conduct several tasks, the actions of terrorists are driven by a hierarchical background. However, in several cases, this hierarchy is hidden and builds up in a self-organized way. For traditional observation techniques, this organization seems to be wide spread, unstructured and loosely connected. Here comes SNA into an important role: collecting small

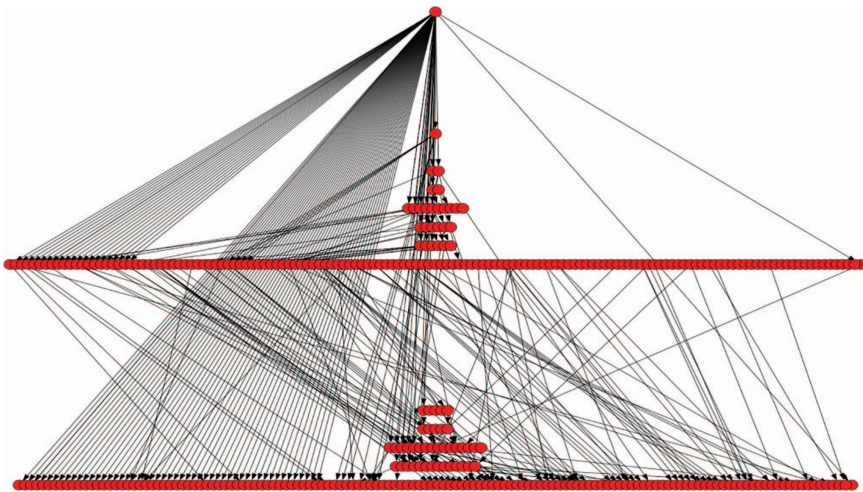



Figure 9.4 Illustration of a possible output of the SNA tool.

pieces of information from huge amount of data results in a holistic picture, where – if data allows it – the unseen hierarchical skeleton can be revealed.

Here, algorithms are implemented for revealing hierarchical structures from flat dataset. New networks are constructed from input data: either from co-occurrence statistics or from directed networks containing loops. Furthermore, quantitative measures are calculated for characterizing the similarity of any network to an ideal hierarchical structure [12].

The upper drawing in Figure 9.4 presents a typical thread-network layout of a forum in the Darkweb. The thread IDs are shown within nodes and

the size of the nodes is proportional to the edges belonging to the given node. Node colours indicate topic groups; links are coloured by the dominant neighbouring node. The lower drawing in Figure 9.4 shows the hierarchical structure of commenters of a Darkweb forum.

9.5 Data Anonymization Tool

We live in an era of technology, where smart devices surround us in all realms of life. These devices feed on our information to generate smart options for us, which at the end help us in making smart decisions. The data gathered by these devices can contain vital personal information such as name, age, location and interest. Alongside these smart devices, nowadays, we tend to rely on social network to broaden the scope of our social interactions. We share personal information such as name and age, we highlight the key things happening in our lives such as places visited, accidents and achievement, we also like sharing our beliefs and interests. Unlike smart devices where adversaries need to corroborate with others to gather information about a single individual, social network data is a source of detailed insight into one's life, thus becoming a bigger threat compared to a single smart device. To mitigate the potential risks, the General Data Protection Regulation (GDPR) was introduced. This new regulation limits the way in which personal data is processed. It limits the ways in which data processing can be done by providing only six lawful ways: Consent, Contract, Legal obligation, vital interest, Public task and legitimate interest⁶. In light of this new regulation, processing social network data becomes tricky. Data collectors, which own the social networks can process this information but after explicitly informing the data owner. This limits the flexibility that third party organizations had. Under the GDPR, all third party companies, who do not have prior consent, need to rely on anonymized data only.

Data anonymization has been around for a while now. It is a process of carefully categorizing social network data into different streams, where each stream undergoes a certain set of tasks. Social network can be divided into three main streams, personal identifiers, quasi-identifiers and non-personal data. Personal identifiers refer to all such parameters that can help identify an individual directly from a large dataset. This mainly constitutes of name, unique-id, contact number and email address. To ensure data anonymization, all such data is removed from the dataset before further processing, thus

⁶<https://ico.org.uk/for-organisations/guide-to-the-general-data-protection-regulation-gdpr/lawful-basis-for-processing/>

reducing the probability of identification of an individual from a large dataset. This probability is further reduced by processing the quasi-identifiers, which on their own have limited meaning but when combined with other quasi-identifiers, can lead to privacy violation. For instance, a dataset containing age information would have less meaning, but when combined with location information would help adversaries in narrowing down their search for an individual and the more quasi-identifiers one has the higher the probability of identification. Therefore, quasi-identifiers are key parameter that all anonymization techniques need to process. The third stream of data deals with the non-personal data. This is set of information that is not connected to any particular individual and can point to anyone in the dataset, for instance, a Facebook post or a twitter tweet can be made by anyone thus, this is considered as non-personal data.

To introduce data anonymization, data analyst carefully analyses the dataset and then narrow down the anonymization approaches that need to be executed. Social network contains three types of quasi-identifiers: numeric, non-numeric and relational information. This therefore means that three separate streams of data anonymization techniques are combined to get results for social network data. Numeric data can be handled by the well-known differential privacy approach [13], where as non-numeric data is handled by k-anonymity (Sweeney, 2002). The relational information is anonymized using a privacy conscious node-grouping algorithm [14]. This anonymized data ensures that no individual can be identified from the processed social network data.

The anonymization techniques applied in this project work in hiding information about all innocent individuals but it also helps terrorist organizations in hiding behind the covers. This as a result puts extra burden on the SNA, CEP, NLP modules. They adapt to working on anonymized social network data and narrow down the search of terrorist organizations. Once identified, the LEAs need to know the identity of the highlighted individuals. To cater for this need, a de-identification approach is also developed in this project, that takes as input the surrogate id's that are provided by the anonymization technique and provide the true identity of an individual. This de-identification algorithm only exists due to the nature of the project and where one can argue that this would make the anonymization algorithm pseudo in nature, it is key to highlight that the de-identification approach only resides with the LEAs thus limiting any adversary from actually identifying individuals and also complying with the GDPR.

9.6 Data Networked Privacy Tool

Intelligence information can be very tricky at times and the nature of this information limits LEAs located in different geographical location from sharing information. On the contrary social networks have no territorial boundaries and terrorist organization can operate from any possible location, making it harder for LEAs to track and tackle them. To overcome this difficulty, the Red-Alert project is equipped with a novel Inter-LEA search algorithm. It limits and controls the amount of information that LEAs located in different geographical location can share with the use of high end encryption algorithm. Under this approach, as shown in Figure 9.5, LEAs are independent in performing their own search and collecting their own intelligence information, they then are requested to populate a list of names of the individuals identified. The second LEA who is looking for a particular individual can search in the encrypted list and find out if one exists or not. The benefit of using high end encryption techniques is of limiting what else an inquiring LEA can see. As such, the LEA inquiring only sees a response in terms of a YES or a NO, therefore hiding all other names in the database. The search query is made with taking a probability attack into consideration,

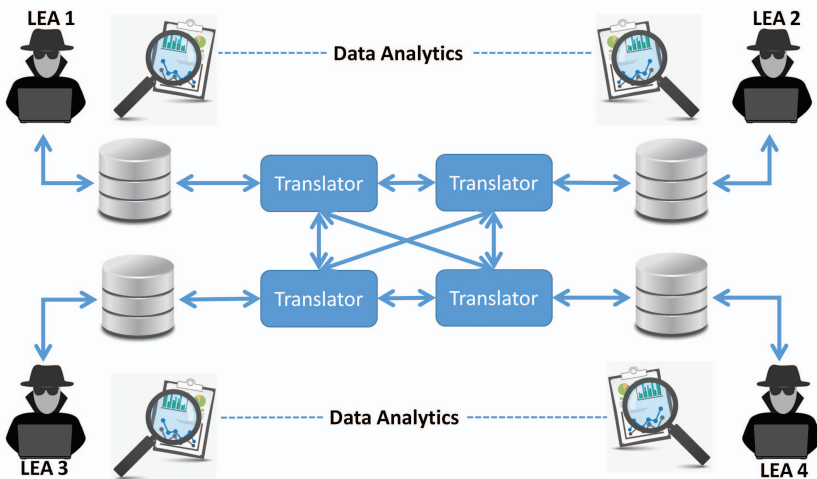


Figure 9.5 Two-layer networked privacy preserving big data analytics model between coalition forces.

thus if an LEA searches for the same name over and over again, there exists no defined pattern. This limits the first LEA (who is hosting the list) from knowing what name is being searched, thus making it a double sided blinded process.

9.7 Integration Component

All the components presented in previous sections will be integrated in one unique solution. The integration component will have therefore some different subcomponents

- **Main System User Interface** provides common look-and-feel to the graphical user interface of the overall RED-Alert System. As shown in Figure 9.6, this component will provide a portal-like user interface for the overall system with common interface placeholders, such as header and footer, main menu, and user interface components hosting through custom common APIs.

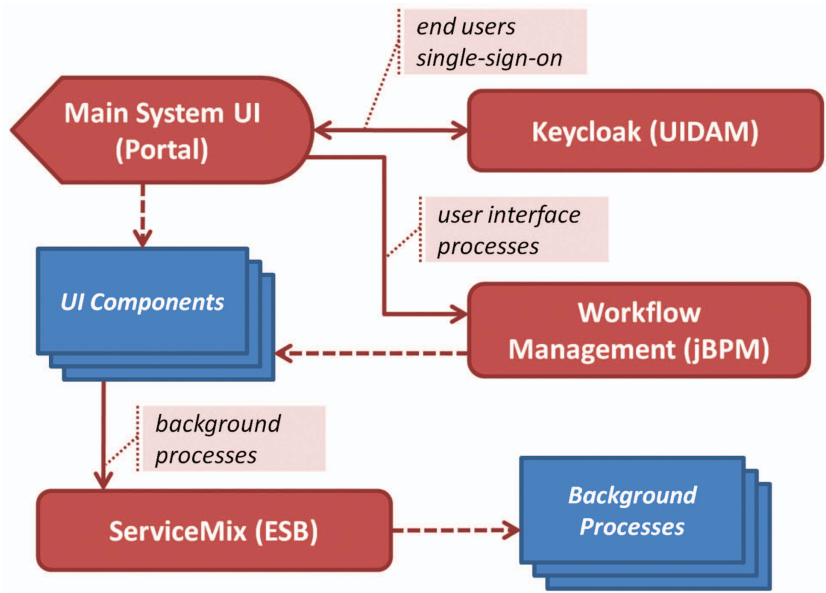


Figure 9.6 Main system user interface – Component interactions diagram.

- **User Identification** and Access Management component will be implemented based on RedHat Keycloak⁷ and will provide the means for identifying users and managing their access to application components, both to front-end user interface and to back-end processes;
- **The Collaborative Workflow/Case Management** component is based on RedHat jBPM,⁸ a light-weight and extensible workflow engine, offering process management features and tools for both business users and developers. RedHat jBPM supports adaptive and dynamic processes that require flexibility to model complex, real-life situations that cannot easily be described using a rigid process;
- **Application Integration Services** component is built with Apache ServiceMix,⁹ an open-source integration container that unifies the functionalities of Apache ActiveMQ, Camel, CXF, and Karaf into a powerful runtime platform you can use to build your own integrations solutions. It provides a complete, flexible, enterprise ready ESB exclusively powered by OSGi;
- **System Interoperability Services** component will be built on top of the Application Integration Services, exposing selected RED-Alert system's functionalities to external system, including existing systems of LEAs;
- **Centralized Audit and Logging** component will be implemented using Audit4j,¹⁰ an open source auditing framework which is a full stack application auditing and logging solution for Java enterprise applications, tested on a common distributions of Linux, Windows and Mac OS, designed to run with minimum configurations, yet providing various options for customization.

Figure 9.7 presents the interactions of the Centralized Audit and Logging component with the Main System UI (Portal), by means of hosting the visual part exposed by the component, and also with the other components of the RED-Alert system, by means of custom common APIs that will allow all components to log entries into a central repository.

⁷<http://www.keycloak.org>

⁸<https://www.jbpm.org>

⁹<http://servicemix.apache.org>

¹⁰<http://audit4j.org>

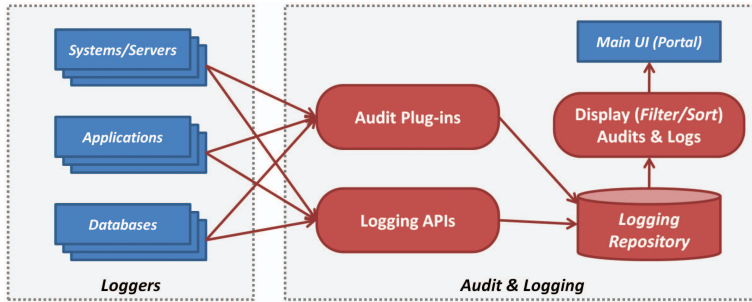


Figure 9.7 Centralized audit and logging – Interactions diagram.

9.8 Future Research Challenges

The next stage of the CEP process is to integrate the recently published standard JSON to finalise the ingestion process (Mongo data source to Kafka topic transfer), and to extend the range of CEP engines to accommodate the SNA (network analysis) data.

One of the major challenges of multimedia extraction is to reduce the number of false positives. We need to make fine grained tuning of SMA tool's components by using larger dataset of a broad range of objects and audio variations. Nowadays data collection, processing and storage have become itself very challenging due to the recently enforced GDPR compliance requirements. The situation is improving with the development of new data management processes and good practices for the data protection. We aim to further improve the performance of SMA Tool and evolve it towards a comprehensive Multimedia Forensics Analysis Toolkit.

Social network analysis is very sensitive to the quality of the available datasets. Further research will aim to develop algorithms for evaluating noisy or biased input datasets. E.g. ensemble averages over possible realizations of networks can shed light on the reliability of predictions.

Another challenge to be addressed is to develop tools for hierarchical visualization of time evolving networks, which helps the analyst in understanding the possible correlations and trends at different scales.

Integration activities will continue with the scheduled iterations towards the piloting phase. These iterations imply adding online streaming capabilities to data acquisition component and expanding the social media channels capabilities beyond Facebook and Twitter, to reach the 7 channels KPI. The Common Schema will be extended with new fields to support these new channels. The security and audit solutions will also be rolled out to all other components to enable the full scope of security requirements of LEAs.

Acknowledgement

This project has received funding from the European Union’s Horizon 2020 Research and Innovation Programme under Grant Agreement No. 740688.

References

- [1] Migration and Home Affairs, “High-Level Commission Expert Group on Radicalisation (HLCEG-R),” Publications Office of the European Union, Luxembourg, 2018.
- [2] OASIS, Privacy by Design Documentation for Software Engineers Version 1.0, 2014.
- [3] N. Marz and J. Warren, Big Data – Principles and best practices of scalable realtime data systems, Manning, 2015.
- [4] Deloitte, Privacy by Design Setting a new standard for privacy certification, Deloitte Design Studio, 2016.
- [5] The Rockefeller Foundation, “Reproduction of Weaver’s memorandum,” 15 07 1949. [Online]. Available: <http://www.mt-archive.info/Weaver-1949.pdf>. [Accessed 12 December 2018].
- [6] D. Kahneman, “Thinking, Fast and Slow by Daniel Kahneman,” *Journal of Social, Evolutionary, and Cultural Psychology*, vol. 2, pp. 253–256, 2012.
- [7] P. Viola and M. Jones, “Rapid Object Detection using a Boosted Cascade of Simple Features,” in *Conference on computer vision and pattern recognition*, 2001.
- [8] S. Ren, K. He, R. Girshick and J. Sun, “Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks,” in *Advances in Neural Information Processing Systems 28 (NIPS 2015)*, 2016.
- [9] I. Goodfellow, Y. Bengio and A. Courville, “Deep Learning,” MIT Press, 2017.
- [10] C. Southall, R. Stables and J. Hockman, “Improving peak-picking using multiple time-step loss functions,” in *Proceedings of the 19th International Society for Music Information Retrieval Conference (ISMIR)*, Birmingham, 2018.
- [11] N. Masuda and R. Lambiotte, A Guide to Temporal Networks, Singapore: World Scientific, 2016.
- [12] E. Mones, L. Vicsek and T. Vicsek, “Hierarchy measure for complex networks,” *PLoS ONE*, 2012.

- [13] C. Dwork, “Differential privacy: A survey of results. In International Conference on Theory and Applications of Models of Computation,” *Springer, Berlin, Heidelberg.*, pp. 1–19, April 2008.
- [14] W. Asif, I. G. Ray, S. Tahir and R. Muttukrishnan, “Privacy-preserving Anonymization with Restricted Search (PARS) on Social Network Data for Criminal Investigations,” 2018.