Multimed Tools Appl
DOI 10.1007/s11042-016-3303-8

Media trustworthiness verification and event assessment through an integrated framework: a case-study

Irene Amerini[1] · Rudy Becarelli[1] · Francesco Brancati[3] · Roberto Caldelli[1,2] · Gabriele Giunta[4] · Massimiliano L. Itria[3]

Received: 28 May 2015 / Revised: 21 January 2016 / Accepted: 26 January 2016

Abstract Nowadays, information is provided through diverse network channels and, above all, its diffusion occurs in an always faster and pervasive manner. Social Media (SM) plays a crucial role in distributing, in an uncontrolled way, news, opinions, media contents and so on, and can basically contribute to spread information that sometimes are untrue and misleading. An integrated assessment of the trustworthiness of the information that is delivered is claimed from different sides: the *Secure!* project strictly fits in such a context. The project has been studying and developing a service oriented infrastructure which, by resorting at diverse technological tools based on image forensics, source reputation analysis, Twitter message trend analysis, web source retrieval and crawling, and so on, provides an integrated event assessment especially regarding crisis management. The aim of this paper is to present

✉ Irene Amerini
irene.amerini@unifi.it

Rudy Becarelli
rudy.becarelli@unifi.it

Francesco Brancati
francesco.brancati@resiltech.it

Roberto Caldelli
roberto.caldelli@unifi.it

Gabriele Giunta
gabriele.giunta@eng.it

Massimiliano L. Itria
massimiliano.itria@resiltech.it

[1] Media Integration and Communication Center, University of Florence, Florence, Italy

[2] National Interuniversity Consortium for Telecommunications (CNIT), Parma, Italy

[3] Resiltech S.R.L., Pontedera (Pisa), Italy

[4] Engineering S.P.A., Roma, Italy

Published online: 03 March 2016

an interesting case-study which demonstrates the potentiality of the developed system to achieve a new integrated knowledge.

Keywords Complex event processing · Crisis management · Image forensics · Trend analysis · Social media · Logo recognition

1 Introduction

The integration of information retrieved from mobile devices, social media and several type of sensors, suggests to exploit the online analysis of a large amount of data allowing to detect and identify dangerous events. Such online approach enables the detection of critical situations as soon as they happen, so that a corresponding reaction can be successfully performed. Many application domains can benefit from this kind of analysis such as surveillance and protection of critical infrastructures and areas, for example: train stations, airports, public squares, world heritage protected areas in some cities of art and so on. The process, starting from the data extraction, leads to the detection of the situation in progress. It introduces several challenges: (i) first of all, it should be highly efficient in order to handle a huge amount of data and detect the situation in progress before it is too late to perform the reaction successfully; (ii) it should be also tolerant to different types of noise, meaning that the process should acknowledge only trusted information from trusted sources, otherwise it could lead to wrong scenario definitions and consequently wrong decisions; and (iii) it should be sufficiently reliable to trust the logged events, including architecture resilience and trustworthy data collection. Complex Event Processing (CEP) [10, 13] systems are widely applied to manage streams of data, in different fields and applications, as business process management, financial services, and also security monitoring, especially for complex, large scale systems where large amounts of information is generated. The *Secure!* Project [1] exactly locates in such an application scenario. The project has studied and developed a service oriented infrastructure which, by resorting at diverse technological tools based on image forensics, source reputation analysis, Twitter message trend analysis, web source retrieval and crawling, and so on, provides an integrated event assessment especially regarding crisis management. This paper presents a case-study in which the *Secure!* Framework has been used to detect critical situation, by managing input data from multimodal sources and providing decision support to the human operators. The rest of the paper is organized as follows: Section 2 introduces some related works while Section 3 presents the *Secure!* Framework logical architecture. Section 4 briefly describes the layers and modules involved in the presented case-study which is instead detailed within Section 6; Section 5 has been dedicated to explain how integration among different modules happens and, finally, in Section 7 conclusions are drawn.

2 Related works

This paper, as already claimed in the Introduction, presents a case-study in which a framework for online trustworthiness verification of social media content and event assessment is practically applied in a real-world scenario. Compared to similar works in terms of approach

[1] Secure! project, http://secure.eng.it/

[6, 8, 12, 19], the paper aims at highlighting several practical complexities in detecting the situation in progress to perform an event assessment on aggregated information. Some works have been done in the past in terms of research papers or international projects to collect, process, and aggregate big streams of social media data and multimedia to discover trends, events, influencers and interesting media contents through the web [6]. Some approaches, in the area of journalism, have been proposed so far but they are mostly still in development; sometimes only some modules of the entire projects are available and often they are not freely available. For example, Verily [20] is a web application designed for the crowd-source verification of information during humanitarian disasters, but it is still in early experimental phase, though it seems to implement a promising approach. Furthermore the crowdsourcing approach based on the "'wisdom of the crowd"', sometimes, is not fully reliable. Instead, it is better to rank evidence from existing information (from Twitter for example) according to the most trusted and credible sources. This is done, for example, in the project Reveal. The objective of the project is to create a trust and credible model able to real-time process evidences by automating news verification steps usually performed by humans and helping the cross-checking tasks. In particular, some works are done on text analysis, extracting and processing fake and genuine claims from tweets referencing suspicious images and videos [19]. Another EU project interested in model, identify, and verify claims through the web is PHEME [7, 24]. In particular, the main task of the project is to identify four types of information (speculation, controversy, misinformation and disinformation) and then to model the spread of such information across social networks and online media to capture human behavior.

3 The *Secure!* framework

The *Secure!* Framework is an integrated ICT framework for enabling crisis and emergency management services. It exploits a process of information retrieval and extraction from several type of heterogeneous and distributed sources, ranging from sensors deployed in the area of interest to web and social networks, in order to detect critical situations and perform the corresponding reaction to eliminate or mitigate the negative effects. The *Secure!* Framework is able to detect critical situations by analyzing events generated from several kinds of analysis over the collected multimedia web resources (i.e. text, images, videos, etc.) and by correlating them with events from other sources (e.g. putting the human-in-the-loop through mobile crowdsensing and crowdsourcing apps) and historical data. For example, a dangerous urban demonstration can be recognized and located correlating data from social network analysis and video analysis. Indeed, the coexistence of a crowd in a precise place (detected by social network analysis and sensor networks deployed in the area) showing a particular symbol (logo) belonging to a dangerous radical group (detected by image/video analysis) could allow the *Secure!* Framework to recognize and locate a dangerous demonstration in progress. The architectural solution adopted for the *Secure!* Framework is based on the most suitable standard for Service Oriented Architecture (SOA). The logical architecture is depicted in Fig. 1.

The architecture is divided into four distinct levels each of which comprises logical components and services. Input data come from the following sources: social media, web sites, mobile devices, sensor networks in critical infrastructures or areas, etc.; starting from the bottom level in the diagram, data are received, collected, homogenized, information are extracted and events are generated, correlated and aggregated in order to produce the *Secure! Situation*. These levels are the following: (i) *Source Data Collector and Media Integration*,

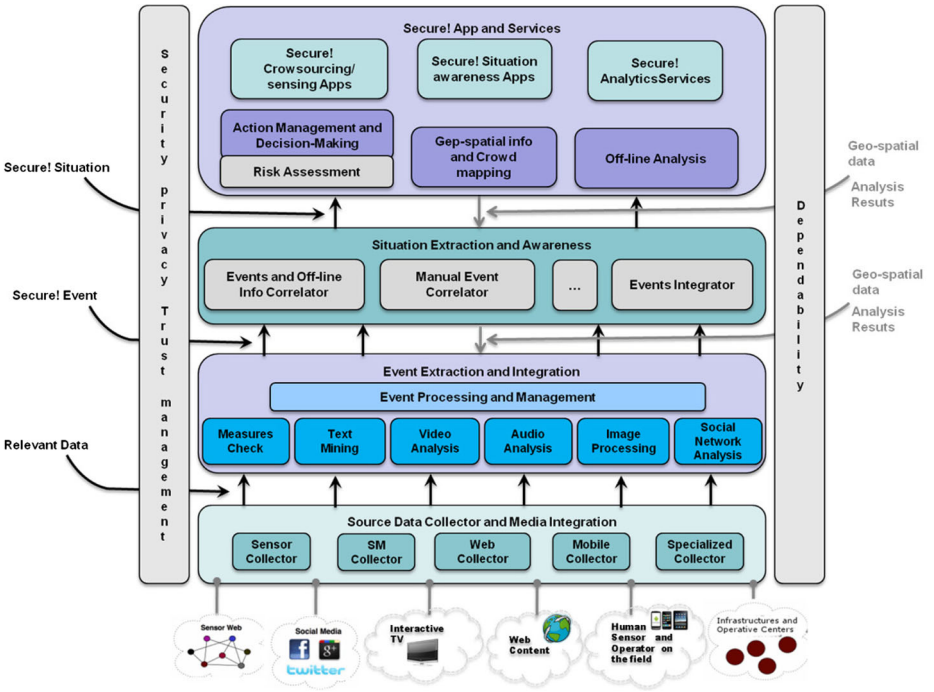

Fig. 1 The *Secure!* framework

which collects heterogeneous data from multimodal sources connected to the system and executes the information extraction of collected data providing relevant data; (ii) *Event Extraction and Integration*, which generates events processing relevant data in input (iii) *Situation Extraction and Awareness*, which processes online events and historical data in order to get the situation awareness and organize the proper reaction; (iv) *Secure!* Apps and Services, which represents the interface between the *Secure!* Framework and the *Secure!* users such as virtual room operators, persons equipped of mobile devices with the *Secure!* Application installed or operators on the field equipped with ad-hoc, safety-critical devices for the situation management. There are also two cross-layers modules for *Security, Privacy and Trust Management* and for *Dependability*.

4 Modules involved within the case-study

This section synthetically introduces some of the modules of the *Secure!* framework that are significantly involved in the case-study presented in this paper.

4.1 The event extraction and integration layer

The *Event Extraction and Integration* layer (Fig. 1: second layer from the bottom) reconstructs the *Secure!* events starting from the relevant data obtained from the underlying *Source Data Collector and Media Integration* layer. First of all, it is responsible for receiving the various types of resources (audio, text, images, videos, measurements, etc.)

extracted from different information sources, integrated into the framework and pre-processed through data cleaning and filtering operations. Then each data is processed by one or more processing blocks on the basis of the kind of resource (two of them specific for text and images will be described in the following subsections). Each of these blocks contains a set of tools that can extract information in terms of features and micro-events. Then these information serves as input to the *Situation Awareness Extraction* layer where the aggregation and the correlation of the *micro-events* will be performed (for a detailed definition of terms related to events see Appendix).

4.1.1 Trend detection component

The aim of the *Trend Detection* component which belongs to *Social Network Analysis* block, is to analyze textual relevant data in order to discover pattern (trend) in the information flow and to identify users who contribute to the propagation of these trends [2, 15]. To this end, text resources are processed by extracting words or phrases (word combinations) that undergo a sudden increase in popularity (bursty keywords) [18], . For example, by analyzing the trend on Twitter it is possible to find out a seismic event after few seconds of its manifestation [21]. If the average frequency of a word, within a time window, is increased, the word is a possible trend; the z-score (standard score) is used to perform the calculation of the number of standard deviations an observation is above a certain mean:

$$z(keyword) = (freq_{current} - average_{historic})/\sigma \tag{1}$$

In particular, the module deals with the identification of trends in textual data provided by Twitter; then the trends are grouped in critical topics through a clustering and finally identifies authors who potentially are organizing, discussing or reporting an interesting event. The module returns a list of emerging trends and the communities (group of users) related to the identified topics.

4.1.2 Logo detection component

This component which belongs to the *Image Processing* block implements a system for automatic recognition of logos and symbols in digital images. Since logos in real images are often occluded by people (i.e. a crowd in a demonstration) or other obstacles, in order to obtain a technique sufficiently robust to partial occlusions and deformations, SIFT descriptors of salient points of an image are used [16]. In particular, the logo L_j is represented by the N_j SIFT feature points (keypoints) detected in the logo image. Detection and retrieval of the logo is performed by comparing the stored local features with those detected within the test image I_i through a matching procedure: a set of matched keypoints M_i is obtained. After that, the localization of the logo in the image is performed by clustering the SIFT matched keypoints selected in the previous step, through the use of a refinement procedure (RANSAC algorithm [11]) that discards outliers points.

4.1.3 Event processing and management component

The *Event Processing and Management* (EPM) component has been designed relying on the Complex Event Processing (CEP) technology [9]. CEP consists of the processing of events generated by the combination of data from multiple sources and aggregated for representing *situations* or part of them. CEP allows an efficient management of the pattern detection process in the huge and dynamic data streams and it is very suitable for recognizing and

correlating events online reducing redundancy, computational complexity and uncertainty. The EPM component recognizes *micro-events*, classifies them depending on the established event taxonomy[1] and produces *complex-events* through information fusion process.

4.2 The situation extraction and awareness layer

Information fusion is intended as the process of merging information from heterogeneous sources with differing conceptual, contextual and typographical representations, reducing redundancy and uncertainty. In *Secure!*, information fusion is one of the enabling process with which situational awareness is achieved through a sequence of operations on events. As soon as the complex-events are produced, they have to be analyzed for checking their information coherence. For this purpose the EPM component interacts with the *Trust Management* component (see Section 4.3) that is able to detect anomalies in the produced complex-events analyzing the spatial localization of the aggregated micro-events exploiting statistical analysis techniques. Finally, complex-events are sent to the *Situation Extraction and Awareness* layer of the framework for building the overall *Secure!* situation used by the decision makers (i.e. operators and domain experts). The *Secure!* situation is built starting from the relevant correlated events at the *Situation Extraction and Awareness* level. *Secure!* operators and domain experts can explore the automatically detected and correlated events by the system and manually interrelate them with other relevant related events. In this way, a semi-automatic event correlation is achieved exploiting the synergy between the machine and the human cognitive abilities.

4.3 Security privacy trust management module

The *Security Privacy Trust Management* module is a cross-layer entity of the *Secure!* Framework. In particular, the *Trust Management* module is active in the following layers: *Source Data Collector and Media Integration* and *Situation Extraction and Awareness*. The task of the *Trust Management* module is to avoid that unreliable data and fake information are input into the system, in order to avert the generation of false negatives and false positives *Secure!* situations. Two main blocks has been designed and implemented for the study of the information accuracy: i) a module to calculate the reputation of the sources (*Source Credibility Evaluator*) and ii) a module that assess the credibility of information (*Resource Credibility Evaluator*). In literature, a well-known mechanism for the calculation of the reputation of a source is the *Beta model* [14, 22] while the computation of credibility of relevant data is made by discriminating among text, images and videos. For the text, the *Resource Credibility Evaluator* is based on a well-known decision technique, adopted in many disciplines (from economics to social sciences to computer engineering), called Analytic Hierarchy Process (AHP) [17]; while for images and videos, the *Resource Credibility Evaluator* resorts at Image Forensic techniques [23] and will be described hereafter.

4.3.1 Resource credibility evaluator for images/videos

The component *Resource Credibility Evaluator* performs an image forensic analysis on digital images (videos) to establish whether these documents are authentic or they have been generated through juxtapositions, retouching, and so on, to alter the meaning of the represented content. Cancellation of a person, replication of an object, juxtaposition of two subjects belonging to two separate contexts are just some immediate examples. The image analysis is performed by resorting at different image forensic tools: each of them is in

charge of looking for diverse kinds of forensics attacks. For instance, the technique which checks for copy-move modifications implements a technique proposed within [3, 4]. Such a method has been chosen because it represents a state of the art instrument in image forensics literature; actually this method has been recently adopted also to detect splicing attack [5]. However, due to the modularity of the *Secure!* platform, it can be envisaged the usage of other tools dealing with image forgery detection.

5 Modules integration

The *Secure!* framework is built on a distributed architecture in which the various modules (services) run on different server machines, often geographically distributed.

Furthermore, such modules presents diverse operational features in terms of input/output resources and, mainly, in terms of computational time; this last necessity has determined that some of them work synchronously but others need to perform asynchronously. Such a specificity reflects on the way the interfacing among modules takes place and, above all, on how different processing events are managed. The dialogue among all the blocks happen via RESTful (REpresentational State Transfer) web services over HTTP and the data are serialized in JSON (JavaScript Object Notation) format. An example of a JSON file, extracted at the output of *Logo Detection Component*, is presented in Fig. 2. The JSON file contains the object *processingResult* which represents the result obtained by applying the specific processing module (in this case the *Logo Detection Component*) onto the *InputImage* whose resource is located at the indicated link. Some properties, such as "credibility" were already valued at "true" by the previous service *Resource Credibility Evaluator* and others, such as "atTime", are instead valued by the process itself. The outputs of the service are provided in *OutputImage* which basically contains the result image as a linked resource, the objects *entities* and *classified* that add additional metadata, specifically, in this case, the description of the group or party connected to the detected logo and the indication of the service for logo recognition.

6 The case study

In this Section, one of the various case studies which have been taken into account by the *Secure!* project to test and evaluate the effectiveness of the proposed architecture is presented to demonstrate how the whole procedure flows. The considered circumstance concerns a political opposition manifestation which has been held in Rome on February 28^{th} 2015 and organized by the right party of *Lega-Nord* together with the movement *CasaPound* against the Italian Government.[2] Furthermore, the situation was potentially high risky also because another manifestation of opposite political extraction was supposed to take place at the same time. The *Secure!* architecture effectiveness has been tested by gathering different information coming from various media sources trying to understand if the system was able to make a reliable assessment of the happening event. In particular, more significant results have been obtained by analyzing what happened on the social network *Twitter* in a defined time period and correlating them with the digital photos appeared both on the web (e.g. web sites of italian newspapers) and attached to the selected tweets. In this case, the

[2] www.repubblica.it/politica/2015/02/28/news/lega_-108382515/

```
{
" processingResult " : [ {
  "id" : 0,
  "atTime" : 1446120103129,
  "circa" : 1446120103129,
  "latitude" : 0.0,
  "longitude" : 0.0,
  "severity" : "medium",
  "eventTag" : [ ],
  "credibility" : true,
  "eventResource" : [ {
    "InputImage" : {
      "id" : 0,
      "title" : "casapound1.jpg",
      ......
      ......
      "imageName" : "casapound1.jpg",
      "link" : "http://yournewspaper.com/LocalNews/casapound1.jpg",
      "rawData" : null
    }
  }, {
    "OutputImage" : {
      "id" : 0,
      "title" : "casapound1.jpg",
      .......
      ........
"credibility" : true,
      "resourceStatus" : null,
      "source" : [ ],
      "resourceCategories" : [ ],
      "imageName" : "casapound1_LOGORECOGNITION.jpg",
      "link" :
"http://<serviceIP>:8080/Logodetectioncomponent/resources/casapound1_LOGORECOGNITION.jpg",
      "rawData" : null
    }
  } ],
  "entities" : [ {
    "Group" : {
      "id" : 0,
      "entityTag" : [ ],
      "roles" : [ ],
      "nameGroup" : "casapound",
      "descriptionGroup" : "Movimento politico italiano. Originariamente CasaPound
nacque in qualità  di primo centro sociale di ispirazione fascista."
    }
  } ],
  "classified" : [ {
    "id" : 0,
    "name" : "Riconoscimento Loghi",
    "uri" : "http://secure.eng.it/ontologySecure/microEvents.owl#Logo_Recognition",
    "description" : "Riconoscimento di loghi rappresentati su striscioni, bandiere,
magliette, ecc"
  } ]
} ]
}
```


Fig. 2 JSON serialized data file

event extraction procedure has mainly involved the components *Social Network Analysis* and *Image Processing* for tweets and image content check-out respectively. The *Twitter* analysis, based on trend identification through semantic text classification, has permitted to collect 5968 tweets on February 28^{th} 2015 which have been successively searched for the presence of pre-defined keywords and hashtags crucial for the *Secure!* project application scenarios. For instance, the hashtag #RENZIACASA(#RENZIGOHOME), among others, has been used, in this phase, to collect all the tweets related to the manifestation. The whole bunch of harvested tweets have been sent as input to the *Twitter Event Producer* (TEP) component which is in charge of determining if a specific event, according to a pre-defined set of category events, is detected or not. The TEP component performs the grammatical analysis of tweet texts and produces a set of classified words (verbs, nouns, dates and places belonging to *Secure!* domain). These words are used to: i) detect the event category, ii) locate the event in space and time (finding relevant dates and places in tweets). In this case 31 events have been recognized as belonging to 7 categories whose histogram is pictured in Fig. 3; it can be pointed out that, by referring to the ground truth, most of the detected events are actually individuated as correct (red columns in Fig. 3).

On the other side, digital photos appeared on the web or attached to some of the selected tweets associated to the *Manifestation* and *Crowding* clusters have been downloaded and processed by two analysis tools of the *Secure!* Framework. It is important to point out that the images under analysis are associated to trusted tweets and reliable web sites. It could be imagined that such images could also come from the *Secure!* crowdsourcing mobile applications acquired by users registered to the *Secure!* system. First of all, the images have to overcome the authenticity check performed by means of the *Resource Credibility Evaluator* tool which is contained within the *Security Privacy Trust Management* module. The *Resource Credibility Evaluator* tool is equipped with image forensic instruments which looks for possible manipulations and anomalies throughout the image itself: if this is the case, such an image is discarded (Fig. 4).

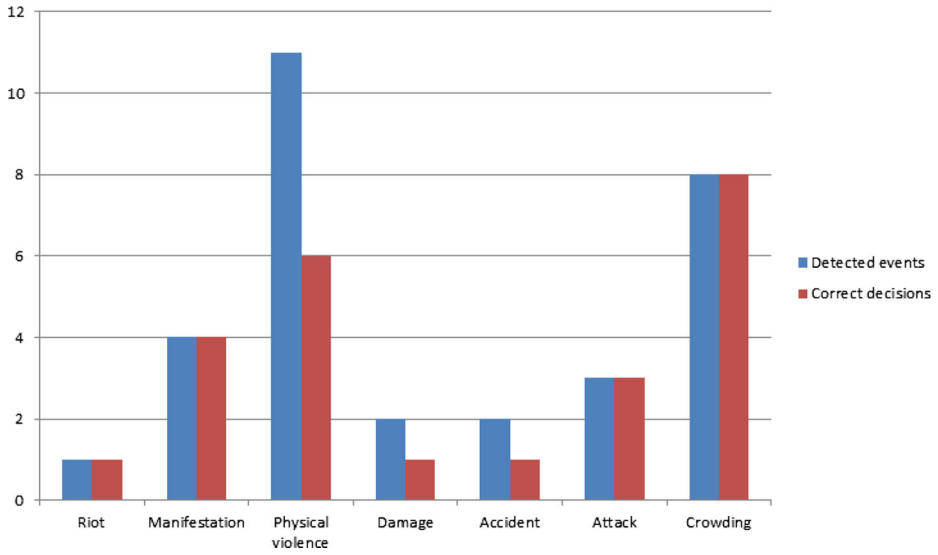

Fig. 3 Histogram of the detected event categories on February 28^{th} 2015. Detected events (*blue columns*) and correct ones (*red columns*)

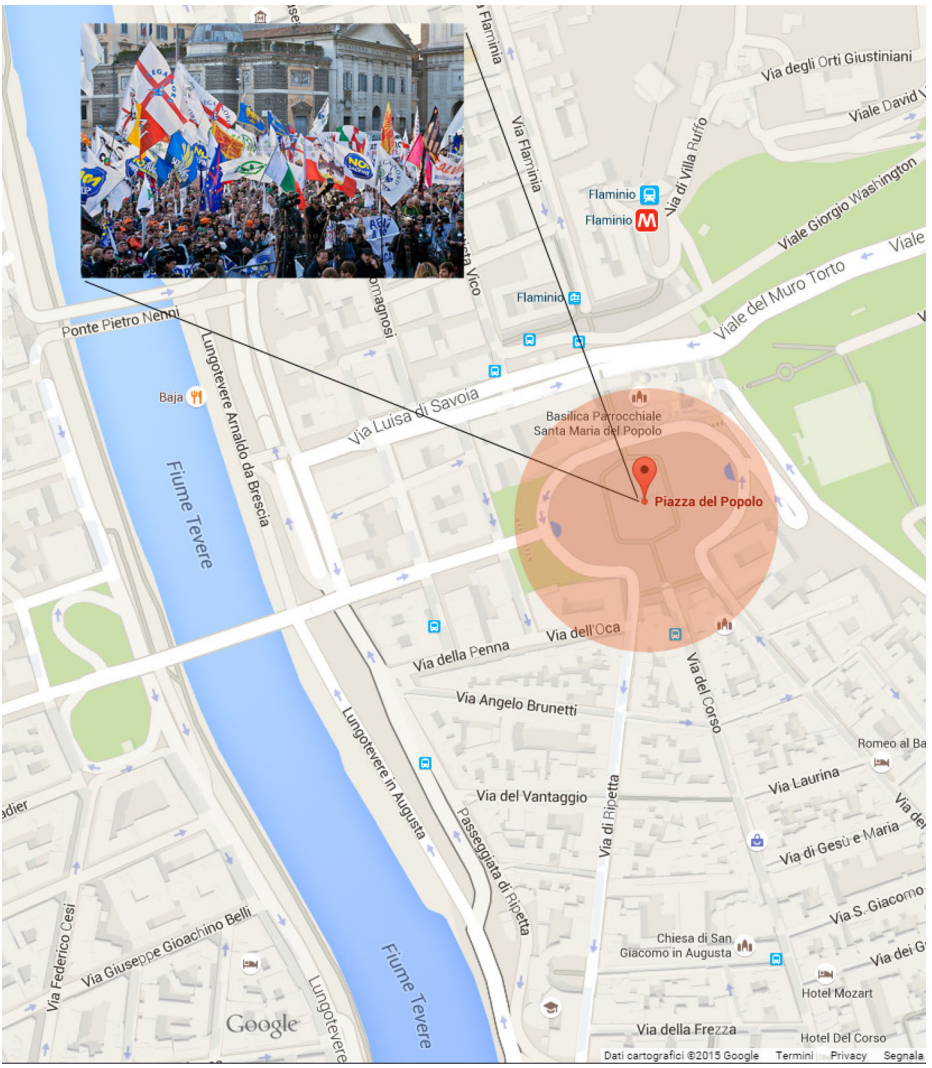

Fig. 4 GPS localization of the manifestation

If the image is validated, then its content is investigated by the other tools of the *Secure!* Framework; in particular, in this case, the *Logo Recognition* tool (contained within the *Image Processing* component) has been involved to reveal the presence of symbols having as reference a pre-defined database containing logos (different versions of them) pertinent with the *Secure!* application scenarios (e.g. political movements, extremist associations, etc.). In Fig. 5, an example of a selected tweet with an attached image containing a logo together with the detection result is pictured. Some of the images are then geo-localized around *Piazza del Popolo, Roma* where the manifestation had actually taken place (see Fig. 4).

In Fig. 6, some of the results obtained by the *Logo Recognition* tool on the processed images are presented. The output of the tool clearly highlights that in some of the images the logos of *Lega Nord* and *CasaPound* are well identified as expected. The logo of the

Fig. 5 A trusted tweet of the *Manifestation* cluster with attached images (*top*) and detected logo (*bottom*)

Fig. 6 Detected logos of political movements and parties

movements supporting the *Lega Nord* leader Matteo Salvini (*Noi con Salvini*) and *CasaPound* such as *Sovranità* are detected as well (in the bottom of Fig. 6). Finally, the whole system has been able to create what is defined as the *Secure! Situational Picture* by integrating some complex events, both derived from *Twitter* and eventually generated from crowd-mobile apps, with the identification of the kinds of involved movements, derived from images appeared on the web and on social networks.

7 Conclusions

The paper has introduced the logical architecture developed within the *Secure!* project and some of its main component modules have been described; specific focus have been dedicated to the integration issues of the diverse implemented functionality. So one of the main goals of the project is achieved, i.e. creating an innovative platform by defining integration mechanisms among diverse technological tools; most of them are state of the art tools but others are still prototypes, though promising and well-performing, obtained from basic research activities. Furthermore a solution regarding how to effectively manage large amount of different data (coming from Internet, sensors and so on) and identify new ways to analyze different kinds of unstructured information is given. In fact, the challenge with Secure! project was related to the unstructured nature of the information in input to the system, because this issue makes difficult to categorize, model and map the data when it is captured and stored. The problem is made worst by the fact that the data normally comes from external sources, whose reliability is often extremely complicated to be confirmed. Therefore these issues are tackled by the proposed infrastructure that is especially suited to analyze web browsing patterns, tweets and transit movements, to predict behaviour and to extract additional hidden information to support activities such as event assessment and critical situation prevision. This is done paying specific attention to the reliability of the data and giving a measure of trustworthiness of the produced answer of the system. One significant case-study, related to the manifestation happened on February 28^{th} 2015 in Rome, has been presented to demonstrate how the whole procedure takes place and to show the potentiality of the entire infrastructure.

Acknowledgments This work was partially supported by the SECURE! Project, funded by the POR CreO FESR 2007–2013 programme of the Tuscany Region (Italy).

Appendix: Definition of terms

The term *event* is defined as "an occurrence within a particular system or domain; it is something that has happened, or is contemplated as having happened in that domain" [10]. In the *Secure!* project this definition considers those events that happen in the real world and are represented in computing systems through structured information. Hence, in the *Secure!* project, each event contains the texture description of the real event, the time/space (when/where it happened), the entity involved and the source that generated it. For sake of clarity we define the terms *micro-event*, *complex-event* and *situation*. The term *micro-event* refers to a simple real event involving one entity only (e.g., people, fire presence, logo recognition, weapon detection) that could be critical or not, therefore the framework needs to analyze it in detail by using other available information. On the other hand, *complex-events* are the aggregation, correlation and integration result of the information contained in a set

of *micro-events* which are correlated by spatial, temporal and causal relations defined by correlation rules. A *complex-event* suggests a *situation* in progress or a part of it (e.g., people demonstration with the presence of crowd and police, vandalism smearing monuments). In the *Secure!* project *complex-events* have been classified through an event taxonomy[1]. With the term *situation*, as defined in [1], we intend "one or more *complex-event* occurrence that might require a reaction". When a critical situation happens a number of specific *complex-events* occur, the commixture and the correlation of them identifies the specific *situation* in progress requiring appropriate reactions, for example providing first aid or police intervention.

References

1. Adi A, Etzion O (2004) Amit - the situation manager. VLDB J 13(2):177–203. doi:10.1007/s00778-003-0108-y
2. Aiello L, Petkos G, Martin C, Corney D, Papadopoulos S, Skraba R, Goker A, Kompatsiaris I, Jaimes A (2013) Sensing trending topics in twitter. IEEE Trans Multimedia 15(6):1268–1282. doi:10.1109/TMM.2013.2265080
3. Amerini I, Ballan L, Caldelli R, Del Bimbo A, Serra G (2011) A SIFT-based forensic method for copy move attack detection and transformation recovery. IEEE Trans Inf Forensics Secur 6(3):1099 –1110
4. Amerini I, Ballan L, Caldelli R, Bimbo AD, Tongo LD, Serra G (2013) Copy-move forgery detection and localization by means of robust clustering with j-linkage. Signal Process Image Commun 28(6):659–669
5. Amerini I, Becarelli R, Caldelli R, Casini M (2015) A feature-based forensic procedure for splicing forgeries detection Mathematical Problems in Engineering 2015. doi:10.1155/2015/653164
6. Boididou C, Papadopoulos S, Kompatsiaris Y, Schifferes S, Newman N (2014) Challenges of computational verification in social multimedia. In: Proceedings of the 23rd international conference on world wide web, WWW '14 Companion, pp 743–748. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, Switzerland. doi:10.1145/2567948.2579323
7. Derczynski L, Bontcheva K (2014) Pheme: veracity in digital social networks. In: Posters, demos, late-breaking results and workshop proceedings of the 22nd conference on user modeling, adaptation, and personalization co-located with the 22nd conference on user modeling, adaptation, and personalization (UMAP2014), Aalborg, Denmark, 7–11 July 2014
8. Derczynski L, Maynard D, Rizzo G, van Erp M, Gorrell G, Troncy R, Petrak J, Bontcheva K (2015) Analysis of named entity recognition and linking for tweets. Inf Process Manag 51(2):32–49. doi:10.1016/j.ipm.2014.10.006. http://www.sciencedirect.com/science/article/pii/S0306457314001034
9. Esper T, EsperTech I (2014) Esper Reference version 4.9.0. http://esper.codehaus.org
10. Etzion O, Niblett P (2011) Event processing in action. MANNING
11. Fischler MA, Bolles RC (1981) Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. Commun ACM 24(6):381–395. doi:10.1145/358669.358692
12. Gupta A, Lamba H, Kumaraguru P, Joshi A (2013) Faking sandy: characterizing and identifying fake images on twitter during hurricane sandy. In: Proceedings of the 22nd international conference on world wide web, WWW '13 companion, pp 729–736. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, Switzerland
13. Itria ML, Ceccarelli AD (2014) A complex event processing approach for crisis-management systems. In: EDCC workshop big4CIP
14. Jøsang A, Roslan I (2002) The beta reputation system. In: Proceedings of the 15th bled electronic commerce conference
15. Kumar S, Morstatter F, Liu H (2014) Twitter data analytics. Springer, Berlin Heidelberg New York
16. Lowe DG (2004) Distinctive image features from scale-invariant keypoints. Int J Comput Vis 60(2):91–110
17. Mao Z, Li N, Winsborough W (2006) Distributed credential chain discovery in trust management with parameterized roles and constraints, vol 4307, pp 159–173
18. Mathioudakis M, Koudas N (2010) Twittermonitor: trend detection over the twitter stream. In: Proceedings of the 2010 ACM SIGMOD intern. conference on management of data, New York, pp 1155–1158. doi:10.1145/1807167.1807306
19. Middleton SE (2015) Extracting attributed verification and debunking reports from social media: mediaeval-2015 trust and credibility analysis of image and video. In: Working notes proceedings of

the MediaEval 2015 workshop, Wurzen, Germany, September 14–15, CEUR-WS.org, ISSN 1613-0073. http://ceur-ws.org/Vol-1436/Paper_43.pdf
20. Popoola A, Krasnoshtan D, Toth AP, Naroditskiy V, Castillo C, Meier P, Rahwan I (2013) Information verification during natural disasters. In: Carr L, Laender AHF, Lscio BF, King I, Fontoura M, Vrandecic D, Aroyo L, de Oliveira JPM, Lima F, Wilde E (eds) WWW (Companion Volume). International World Wide Web Conferences Steering Committee / ACM, pp 1029–1032
21. Sakaki T, Okazaki M, Matsuo Y (2010) Earthquake shakes twitter users: real-time event detection by social sensors. In: Proceedings of the 19th international conference on world wide web, WWW '10. ACM, New York, pp 851–860. doi:10.1145/1772690.1772777
22. Sherchan W, Nepal S, Paris C (2013) A survey of trust in social networks. ACM Comput Surv 45(4):47:1–47:33. doi:10.1145/2501654.2501661
23. Stamm M, Min W, Liu K (2013) Information forensics: an overview of the first decade. Access, IEEE 1:167–200. doi:10.1109/ACCESS.2013.2260814
24. Zubiaga A, Liakata M, Procter RN, Bontcheva K, Tolmie P (2015) Towards detecting rumours in social media. In: AAAI workshop on AI for cities

Irene Amerini received the Laurea degree in computer engineering and the Ph.D. degree in computer engineering, multimedia, and telecommunication from the University of Florence, Florence, Italy, in 2006 and 2010, respectively. She is currently a Post-Doctoral Researcher with the Image and Communication Laboratory, Media Integration and Communication Center, University of Florence. She was a Visiting Scholar with Binghamton University, Binghamton, NY, USA, in 2010. Her main research interests focus on multimedia content security technologies, secure media, and digital and multimedia forensics.

Rudy Becarelli He has been involved in research and development activities at University of Florence since 2004, currently he is pursuing his PhD in Computer Science, Systems and Telecommunications. Research activities mostly concern with digital watermarking, interactive TV and image forensics. He is expert in design and development of J2EE applications for data exchange and marshalling with rich client platforms. He has also experienced OAI-PMH server development, digital library customization and profiling, MHP application development and optimization.

Francesco Brancati tooks his Master degree in Computer Science at the University of Firenze in 2008 and his PhD degree in Computer Science at the Resilient Computing Lab - University of Firenze in 2012. His research activity mainly focused on adaptive and safe estimation of different sources of uncertainty to improve dependability of highly dynamic systems through online monitoring analysis. Currently he is coordinating the R&D projects and activities in Resiltech.

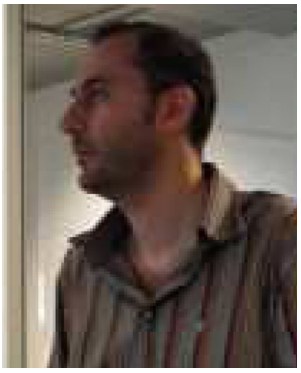

Roberto Caldelli received the degree in electronic engineering and the Ph.D. degree in computer science and telecommunication from the University of Florence, Florence, Italy, in 1997 and 2001, respectively. From 2005 to 2013, he was an Assistant Professor with the Media Integration and Communication Center, University of Florence. In 2014, he joined the National Inter- University Consortium for Telecommunications (CNIT) as a Permanent Researcher. His main research activities, witnessed by several publications, include digital image processing, interactive television, image and video digital watermarking, and multimedia forensics. He holds two patents in the field of content security and multimedia interaction.

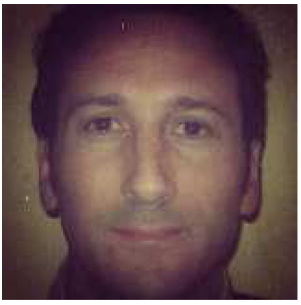

Gabriele Giunta Project Manager at R&D Laboratory of Engineering Ingegneria Informatica S.p.A. Specialties: Intelligent Transport System technologies, Semantic Web technologies, HCI, Business Process Modeling tools, Business Process Execution languages.

Massimiliano L. Itria Software Engineer at ResilTech S.r.l. Specialities: Software Design, Semantic technologies, Compmlex Event Processing technologies, Software Monitoring Systems, Stochastic Modelling.