Newsletter No. 1 July 2012

Search engine for multimedia environment generated content

Welcome to SMART First Newsletter

Welcome to the first edition of the Newsletter of the SMART project (www.smartfp7.eu), which is intended to provide information on the progress of the project's work (including the main current technical achievements), as well as information on its interaction with stakeholders (including other projects, the media and the open source community). The SMART newsletter will be published approximately each semester, depending on the evolution of the project's work and the accomplishment of relevant milestones. On behalf of the SMART consortium, I hope that you will enjoy reading the SMART newsletters and that you will find them useful.

John Paul Moore, Atos, SMART Project Coordinator.

Objectives of SMART

The main goal of the SMART project is to research and implement a scalable open source next generation multimedia search engine, which will be able to search information stemming from the physical world. The SMART multimedia search engine will be able to answer queries based on the intelligent collection and combination of sensor generated multimedia data, based on sensors and sensor processing algorithms that match the context of the query at hand. The matching of the queries with sensors and the sensors' processing algorithms (notably audio and video processing algorithms) will be based on the sensors' context and metadata (e.g., location, state, capabilities), as well as on the dynamic context of the physical world. Events of interest are extracted from the audio/visual streams of the various sensors using various appropriate multimedia processing algorithms (such as face detectors, person trackers, classifiers of acoustic events, crowd analysis components and more). Furthermore, SMART will be able to leverage social network information in order to facilitate the answering of social queries over a multitude of sensor data.

The main objectives pursued by SMART are:

1. To provide an open scalable and extensible framework for searching (multimedia) information on content stemming from a multitude of sensors.
2. To develop a set of robust and operational algorithms in a dynamic far-field environment for audio and visual processing such as acoustic event classification and crowd analysis.
3. To implement schemes for the dynamic orchestration/ organization of societies of sensors (i.e. cameras, microphones) and algorithms (e.g., acoustic processing, visual scene analysis), to respond and anticipate search queries.
4. To research technologies for the combination of information derived from sensors with information stemming from social networks, with a view to delivering relevant context-aware social information along with the environment information context.
5. To research utility metrics, incentive schemes and business models, encouraging stakeholders (i.e. sensor infrastructure providers, technology providers) to contribute sensors and algorithms to the SMART engine, and encouraging end-users to use the search engine.

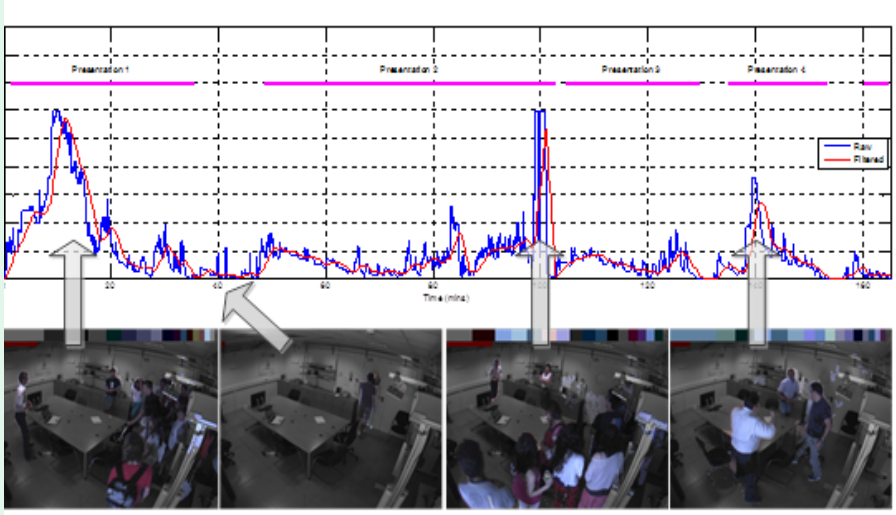

Figure 1. Crowd analysis in SMART.

Crowd Analysis

Crowds are a main source of events for SMART. At a first level of analysis, we are interested in quantifying their density, which is based on an adaptive foreground segmentation. This segmentation is based on a variant of the Stauffer's algorithm. Each pixel in the image is modeled as a Gaussian mixture, which is updated using a spatio-temporally adapted learning rate. The model is then utilized to decide if the pixel belongs to the foreground by adaptively thresholding the accumulated sorted weight of the Gaussians. The foreground mask is then cleaned-up by a shadow removal and a morphological clean-up. The weighted density of foreground pixels

Iadh Ounis – Antonis Litke – Aristodemos Pnevmatikakis – Daniel Field – John Soltados – John Paul Moore – Regis Riveret – Tomas Garcia Fresno – Zvi Kons

gives the crowd density a continuous metadata stream. Using this stream, the low-level events on crowd appearance or disappearance are easily obtained via thresholding.

An example is illustrated in Fig. 1, which shows the crowd density metric as a function of time. This has three large peaks, for which the corresponding processed frames indicating the foreground blobs are given. A fourth frame is an example of minor activity. The extracted low level metadata indicating crowded intervals are also shown as labeled horizontal lines.

SMART will also be extracting other crowd-related metadata such as motion patterns (crowd flow direction, crowd convergence and crowd divergence) and colour (colour trends of crowds).

Figure 2. SMART sensors location in Santander.

Audio & Speech Processing

Audio sensors will have a significant part in the SMART system's ability to access the physical world. This will be done by processing the audio streams and identifying different events by the sounds that they generate.

In SMART, we are currently working on the creation of tools for the classification of continuous audio events (such as crowd noise and music). Those tools operate on the audio stream by first segmenting them into short, fixed length segments and then generating a representation of the audio within the segment using various features. The classification of the extracted features into the different event types is done using a multilayer perceptron networks. The training of those networks employs Deep Belief Networks (DBN) techniques, which have recently proven to be very useful for training speech classifiers in phoneme recognition tasks.

Information collected from the users is another important input of the SMART system. This is usually done with textual interfaces (such as Twitter). In SMART, we take another step forward by allowing the users to provide us with voice messages. We are currently working on two aspects of this technology:

Speech transcription: Speech recordings and textual data will be used for training models for the speech transcription engine.

Speaker verification: some of the speech recordings will also be used to train and test the speaker verification engine. At the same tim SMART is working to improve the speaker verification algorithms and to add new ones into the verification engine. Recently a light JFA (Joint Factor Analysis) module has been added to the speaker verification engine. This module improves the verification rate to be almost as good as the classical JFA methods, while requiring much less computational resources.

SMART Sensor Locations

The locations depicted in Fig. 2, showing a map of Santander, have been pre-selected based on (i) their proximity to the city centre implying a huge affluence of people and, (ii) their access to the Santander Council's optical fibre network. The pre-selected locations are:

1. Plaza del Ayuntamiento: This busy town square is the geographical and business heart of the city.
2. Plaza del Príncipe: This is one of the most crowded places in the city, with a huge affluence of people from commercial activities.
3. Mercado de la Esperanza: In this square, just behind the Town Hall, is located one of the most popular markets in the city.
4. Alameda de Oviedo / Plaza de Numancia: This other location is also one of the most popular venues in the city of Santander.
5. Paseo de Pereda: Pereda's Park is one of the most emblematic and popular locations of the city of Santander.
6. Mercado de México: Mexico's Square is next to Mexico's Market and the local bullring, where the flow of people and vehicles is more than assured throughout the year.

The number of possible collection sites is limited by the number of available video cameras and microphones. The final decision for the data collection locations will depend on the findings from the data pre-collection stage.

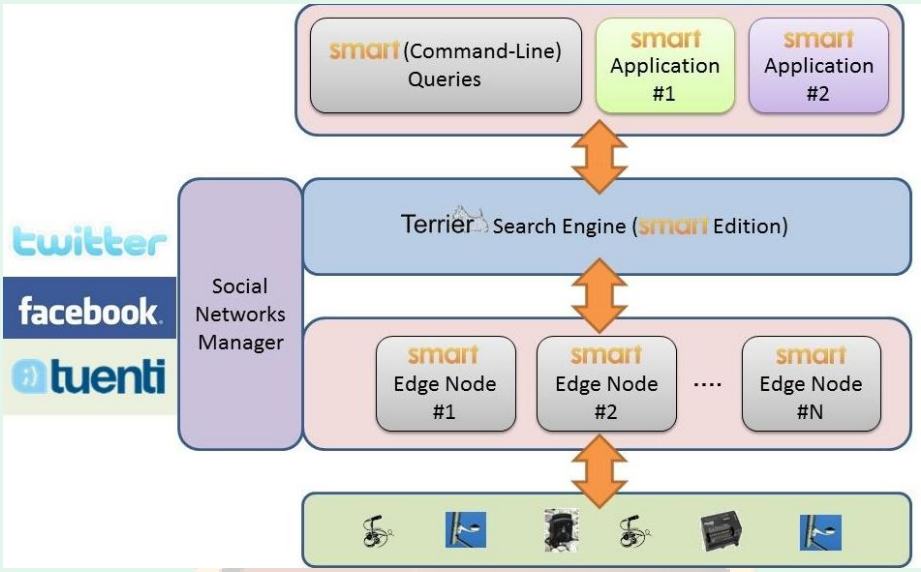

Figure 3. SMART Architecture.

Ethical Issues in SMART

SMART entails the acquisition and processing of content and context from the surrounding environment, including outdoor urban environments. To this end, the SMART applications will deploy perceptive algorithms enabling context extractions from the physical environment/world. Such components involve ethical issues (and in particular privacy issues) as information could be captured relating to individuals in the vicinity of sensors. Although this is outwith the scope of the project itself, the open nature of the architecture could theoretically enable third parties to develop algorithms that track people and their behavior. This is a potential concern to stakeholders and any such developers must be aware of what is permissible or prohibited by law.

In order to ensure compliance with applicable laws and regulations by the project and its participants, and to guide any third party developers, the SMART project is investigating the ethical and privacy implications of its technologies, including relevant measures that can be taken in order to alleviate privacy concerns. As a first step, a survey of applicable laws and regulations, including the citizen's rights, has been conducted and the relevant areas have been analyzed in depth. These revolve principally around the rights of free movement and free speech and the right of data protection. Actors using the SMART technology must ensure compliance with the data protection procedure (e.g. informed consent, secure storage and processing of data and the various requirements on correcting or removing personal information), especially when that data is classifiedd as special or sensitive. In light of this ethics analysis, the SMART project is ensuring that the sensors deployed in the project are incapable of collecting unauthorized personal data, as well as ensuring procedural correctness for the collection, processing and storage of authorized data.

As a second step, the project has already established collaboration links with Data Protection Authorities (DPAs) in Spain and Greece. DPAs are a crucial part of the EU's practices towards data protection, since each DPA serves as the single point for requesting approval and advice about ethical and privacy issues (at national and regional level), including the protection of personal data. This collaboration consists of consulting with the DPAs on the planned deployment of the technology within the project and/or receiving explicit approval for the use of contextual sensors. Future third party deployers who intend to use SMART are advised to do the same. As a result, the project and the deployers of SMART sensors and edge nodes (i.e. proxies to the physical world comprising one or more sensors and perceptive components) can ensure the ethical and legal nature of their deployment. At the time of writing, the SMART consortium is awaiting approvals from the Spanish DPA to deploy SMART sensors and perceptual components in the City of Santander. As a contingency for technologies, which cannot be deployed in public spaces, SMART intends to test them in controlled environments using actors who have provided informed consent for their participation in the SMART data collection and experimentation.

Interested readers, and in particular future deployers of the technology, should consult the series of ethical analysis documents, which SMART will produce throughout the project (namely deliverable D7.8). The first version is available and updates will be made on a yearly basis. These deliverables will provide concrete details on the ethical implications of the SMART technologies, while also outlining relevant measures for alleviating privacy concerns.

SMART Architecture Overview

The SMART architecture (see Fig. 3) is layered and comprises the following layers:

Layer 1: The Edge Nodes communicate with the physical world, by collecting, processing and fusing sensor and social network streams.

Layer 2: The Search Engine (based on the Terrier.org search engine) performs scalable search and retrieval over multiple SMART edge nodes. It offers query processing and Smart-Reduce functionalities.

Layer 3: The Applications are built based on queries to the SMART search engine and employ reusable mashup libraries for visualization.

You can access more details about the public version of the SMART architecture at the SMART Web Site (www.smartfp7.eu).

The SMART Open Source Project

SMART is designed as an open source framework, extensible in terms of sensors, multimedia processing components, and event retrieval models.

The main components of the SMART search engine are built upon the Terrier open source information retrieval platform, allowing for the real-time indexing and retrieval of multiple and massive-scale sensor and social networks streams.

The SMART open source framework is designed to benefit from the power of the open source development philosophy, by enabling application developers and organizations to build new tailored services and products on top of the SMART open source infrastructure. SMART will form an open source community for sustaining and evolving its components. It adopts a crowd-sourcing approach to the deployment of physical sensors, social networking feeds and associated repositories, which will become searchable through SMART.

Thanks to an open specification for describing data streams, the SMART open source framework facilitates prospective information providers (including sensor infrastructure providers) to connect egde nodes and data feeds to the SMART search engine. Hence SMART is designed to integrate a variety of community-based sensor feeds contributed by third parties such as smart cities, sensor deployers and individuals. Likewise, the SMART open source infrastructure supports virtual sensors streams such as data feeds stemming from social networks. In this case, SMART allows social sensors (e.g. gender analysis or sentiment analysis filters on Twitter), to be used while developing applications for smart cities.

Finally, SMART will be released under the business friendly MPL 2.0 (Mozilla Public License) in order to facilitate service integrators to build custom search applications in response to specific business requirements. The open source application layer will make it easier for such services to be rapidly implemented. In this way, SMART intends to support both a public crowd-sourcing paradigm and a private enterprise-related one. The first release of SMART is planned before the end of 2012.

SMART Twitter Demo

The SMART Twitter Noun Frequency Demo (developed by SMART partner TELESTO) demonstrates a simple Twitter filtering application, which can be used as data feed to the SMART search engine. The goal of this application is to demonstrate how data stemming from social networks can be utilized within SMART to enable participatory sensing. To use this demo, a keyword or a hashtag must be provided as an argument for a Twitter query. Next, by filtering the words within the retrieved tweets, a noun histogram is created. Most frequent terms are presented as hyperlinks within a histogram or within a 3D spherical cloud. As time goes by, the tags that appear within the cloud for each input keyword will change, depending on the content that Twitter users will post. Try our demo out at:

http://telesto.zapto.org:81/TwitterNounFilter/

SMART in the News & Publications

SMART deals with a number of hot topics and has already received extensive media coverage.

You may have listened to Iadh Ounis being interviewed live on BBC Radio Scotland about the project. The interview is no longer available on the BBC website, but it shall be soon available on the SMART website www.smartfp7.eu.

Other media coverage includes: BBC News: Glasgow University pioneers internet sensory search engine, STV.tv: Computer scientists on brink of inventing 'sensory search engine', Glasgow University News: University of Glasgow developing new type of internet search engine, Scotsman.com: Sensor search engine developed, Techworld.com: Scottish scientists build search engine for 'Internet of Things', ZDnet.co.uk: Glasgow's SMART search engine senses cities, Evening Times: Boffins make a SMART move to rival Google, Herald Scotland: Scots work on new web search engine, IBN Live: New sensor-based search engine to be developed, The Times of India: Scientists developing sensor-based search engine, The Telegraph: New search engine aims to Google the real world, BBC, Technology Section, Featured Article: Researchers work on smart city search engine, TechWeek europe: Smart City Search Engine Uses Sensors: University of Glasgow search engine project will use social media and sensors to take the pulse of a city., Digital Spy: SMART search engine lets users search real world, SmartPlanet: A search engine for smart cities, Irish Independent: New search engine aims to Google the real world, Economics Times: New sensor-based search engine to be developed, The Telegraph: SMART searching, Westend TV: New search engine aims to Google the real world, Techeye: Scientists create real-time smart city search engine: "SMART" provides live local updates, New Electronics: Researchers developing 'smart city' search engine, IT Pro: Researchers outline real-time search engine plans, ITProPortal: University of Glasgow researchers working on new search engine to answer queries Google can't, Wall Street Journal (blog): New Search Engine Combines Twitter with Sensors, The Cutting Edge (Feature Article): Search Engine Aims to Quiz Sensor

Networks, ComputerWorld (Feature Article): Scots develop search engine which uses social media data: A new search engine is in development which will tell users what their friends are up to, futura-sciences.com: Smart, le moteur de recherche qui interroge des capteurs en ville (SMART: the search engine that queries the sensors in town).

Papers published so far include:

Irene Schmidt, John Soldatos and Paul Moore, «Multimedia Search and Retrieval over Integrated Social and Sensor Networks», Third International Conference on Computational Aspects of Social Networks (CASoN 2011), October 19-21, Salamanca, Spain.

John Soldatos, Moez Draief, Craig Macdonald and Iadh Ounis, «Multimedia Search and Retrieval over Integrated Social and Sensor Networks», In the Proceedings of the WWW2012, conference (EU Projects Track), Lyon, France, April 2012.

Other SMART publications are coming: stay tuned in the next issue of our newsletter.

SMART Partners

Atos

IBM

PRISA DIGITAL

Athens Information Technology – Center of Excellence for Research and Education – 2002

AIT CENTER OF EXCELLENCE FOR RESEARCH AND EDUCATION

Perceptum ex Optimus

Telesto Technologies – Information Communication Technologies

S3log

AYUNTAMIENTO DE SANTANDER

Imperial College London

University of Glasgow – Via Veritas Vita