SEVENTH FRAMEWORK PROGRAMME

Networked Media

Specific Targeted Research Project

SMART

(FP7-287583)

Search engine for MultimediA environment generated contenT

D7.6 Open Multimedia Search Best Practices, Guidelines for Policy Development

Due date of deliverable: 31-08-2014

Actual submission date: 15-09-2014

Start date of project: 01-11-2011 Duration: 36 months

Summary of the document

Code:	**D7.6 Open Multimedia Search Best Practices, Guidelines for Policy Development**
Last modification: **State:**	26/08/2014 Final
Participant Partner(s): **Author(s):** **Fragment:**	ATOS, AIT John Soldatos, Aristodemos Pnevmatikakis, Dyaa Albakour, Romain Deveaud, Iadh Ounis, Josemi Garrido No
Audience:	☒ public ☐ restricted ☐ internal
Abstract:	The purpose of this deliverable is to document some of these experiences in the form of best practices for open multimedia search (over environment generated content), which include also lessons learnt from the SMART deployment.
Keywords:	Open multimedia search
References:	DoW

Table of Contents

1 Executive Summary

1.1 Scope

SMART have developed an open source search framework enabling the development and integration of applications that search for information over a variety of sensor and social networks. The framework has been already validated in the scope of integrated applications, which have been developed over sensors deployed in the city of Santander. All the technical components of the SMART architecture have been developed and delivered in multiple iterations, thereby boosting their continuous improvement, but also the acquisition of experiences associated with their deployment and use in search applications. At the same time the real-life deployment of the SMART applications has led the partners to deal with a number of administrative and deployment issues, including ethical and privacy management issues. Overall, on the basis of the above activities the consortium has accumulated experiences associated with the development, deployment and use of search applications over data streams that stem from sensor and social networks. The purpose of this deliverable is to document some of these experiences in the form of best practices for open multimedia search (over environment generated content), which include also lessons learnt from the SMART deployment.

1.2 Audience

This deliverable is destined to third-parties (i.e. outside the SMART consortium), which have an interest in building and/or deploying and/or operating search applications over integrated sensor and social networks. These include:

- Solution developers and integrators building search solutions over social networks and sensor networks, including integrators of the SMART open source community.
- Infrastructure owners and deployers in SMART cities, including public authorities and smart city authorities.
- Providers and operators of search services (including search engines and related search applications) in smart cities.

1.3 Summary

This deliverable presents a range of best practices and lessons learnt from the SMART project activities. The purpose of these best practices is to facilitate solution developers, integrators and infrastructure providers to successfully engage in the development and deployment of social/sensor applications for urban environments using SMART or similar frameworks. The best practices and lessons learnt target the following areas of applications development, deployment and integration:

- The ever important ethical and privacy issues, which are integral parts of any non-trivial deployment in the urban environment.
- The modeling of the different types of data feeds from sensor networks and social networks.
- The interfacing to a wide range of data sources including physical devices, algorithms, data clouds and more.
- The real-time indexing and retrieval of information, which is particularly important in the scope of low-latency applications that leverage high-velocity data.
- The implementation of visual processing and acoustic processing algorithms over cameras and microphones deployed (respectively) in the urban environment.

1.4 Structure

The rest of the deliverable is structured as follows:

- Section 2 illustrates best practices and recommendations associated with ethical and privacy management in the scope of search applications that leverage environment generated context.
- Section 3 illustrates tactics for modelling sensor and other types of data streams in the scope of multimedia search applications over sensor and social networks.
- Section 4 is devoted to a review of the different types of data sources, which can contribute data and information to SMART-like multimedia search applications.
- Section 5 discusses best practices about personalized recommendations and real-time indexing for searching information over social and sensor networks, especially in cases where data streams with high ingestion rates are involved.
- Section 6 is devoted to recommendations about audio processing (notably in the areas of acoustic event identification).
- Section 7 documents recommendations and main lessons learnt from the video processing technologies of the project, in areas such as crowd analytics.
- Section 8 concludes the deliverable.

2 Best Practices and Lessons Learnt on Ethical and Privacy Management

2.1 Sensor Deployment and Data Privacy Authorities

Several SMART technologies (e.g., crowd analysis, acoustic event detection) rely in the deployment of sensors (such as cameras and microphones) in the urban environment. This deployment is subject to approval by national Data Privacy Authorities (DPA) and should not be taken for granted. The relevant approvals should certify that the sensor deployment and the way it will be used will not compromise citizens' privacy. SMART has allocated significant effort towards consulting with DPAs prior to having its sensor deployments approved. As part of this effort the project has taken all necessary measures in order to get approval for its sensor deployment activities, notably in terms of A/V sensors deployment.

2.2 Privacy Implications of Social Networks

SMART has also developed modules for accessing and processing data streams and information from social networks. Information derived from the latter may also raise privacy concerns. The exploitation of social media in the scope of multimedia sensor and social search applications is governed by the privacy policies of the social media platforms involved. This raises privacy issues, especially in cases where information about a user can be publicly associated with other information outside the social network and in cases where personal information is shared without explicit consent by the user.

2.3 Recommendations

Based on the SMART experiences, the following recommendations and guidelines can be providers for future deployers of sensor technologies in the urban environment for research purposes:

- **Recommendation 1.1**: The deployment of cameras and other sensors in public spaces required written approval by the DPAs. It is recommended that you understand the scope and process of this approval prior to applying for it to the DPA. Consider also consulting the DPA on issues relating to the application process e.g., what is should contain, any supplementary material that should be provided, any measures that should be listed and more. Your application should explain the objectives, scope and rationale of the sensor deployment.
- **Recommendation 1.2:** Policy makers should consider «research» (and R&D activities) as one of the purposes that would be legitimate/eligible for deploying sensors in urban environments, following DPA approval. In several laws this is not mentioned explicitly and the final decision is taken by on subjective judgments of the DPA.
- **Recommendation 1.3:** It is recommended that you take measures (including technical measures) for ensuring citizens' privacy. This may for example include measures associated with the placement of cameras (e.g., heights), but also the type of sensor processing applied.
- **Recommendation 1.4:** Pay attention to the privacy implications of the use of data from social networks, stemming from their privacy policies. Special attention should be given in cases of social networks that do not comply with the requirement of explicit informed consent to process user's data.

3 Best Practices and Lessons Learnt on Sensor Data Feeds Modeling and Representation

3.1 Data Feeds Modeling options

Information search and retrieval in SMART is performed over large collections of data streams (stemming from physical and virtual sensors). A unified modeling of the various heterogeneous data streams is a key to enabling their unified processing and search. Unified models emphasize both on syntactic unification (use of a common format) and on semantic unification (definition and use of common semantics). SMART has selected an approach that represents a compromise between the verbose semantically rich technologies for modeling data feeds (e.g., Semantic Sensor Networks ontology from the respective W3C Incubator group) and the extremely simple (but semantically poor) CSV and JSON format for modelling data feeds. This compromise is based on SMART's custom XML/JSON schemas and enables the feeds descriptions to capture context, while ensuring simplicity in their processing.

3.2 Recommendations for Data Streams/Feeds Modeling

The following recommendations on data feeds modeling are provided to potential search solution providers and integrators:

- **Recommendation 2.1**: Large scale search applications benefit from the collection of multiple heterogeneous data streams for physical and virtual sensor sources. As part of this collection the different formats and semantics of the diverse data streams need to be unified.
- **Recommendation 2.2**: Recent semantic web standards provide the means for semantically interoperable processing of different data streams, yet they incur processing overhead and a learning curve. They should be adopted when semantic interoperability is a primary concerned.
- **Recommendation 2.3**: Simple data modeling formats should be used for performance and programming/development simplicity. These can nowadays be combined with the semantic standards based formats (e.g., ontologies) if needed, as recent trends also show (e.g., the emergence of JSON-LD (Javascript Object Notation – Linked Data) format).

4 Best Practices and Lessons Learnt on Types of Data Sources & Feeds for Search Applications

4.1 Types of Data Sources/Feed supported in SMART

SMART has developed a versatile system for the collection and processing of data feeds from sensors and social networks, as well as for the subsequent development of search applications over them. The versatility of the system is reflected on its ability to process and deal with:

- Feeds from physical sensors and wireless sensor networks.
- Feeds from perceptual components (notably signal processing algorithms) operating over sensor streams.
- Feeds from social networks, including Twitter, Foursquare and other popular social network platforms.
- Feeds derived from the LinkedData cloud (such as information contained in DBPedia).
- Feeds inferred based on reasoning over the above-listed data streams.

SMART provides open source components and accompanying documentation of its components that deal with the collection and processing the above data feeds. For example, the perceptual components which operate on the sensor data streams are detailed in GitHub at https://github.com/SmartSearch/Edge-Node/wiki/Perceptual-components, while the social network manager (SNM) that provides services for requesting data from social networks and filters them to derive social metadata is available at: https://github.com/SmartSearch/Edge-Node/wiki/Social-Network-Manager. Also, the linked data manager which collects linked data of interest to a certain point of presence of the SMART system is documented at https://github.com/SmartSearch/Edge-Node/wiki/Linked-Data-Manager.

4.2 Recommendations

Based on the above implementation experiences, the following recommendations for solution provides can be drawn:

- **Recommendation 3.1**: Sophisticated search applications over sensor and social networks can take into account a variety of data streams from both virtual and physical sensors. These data streams can be technically handled (e.g., modeled, programmed, processed, managed) in the same way regardless of the type of data feeds.
- **Recommendation 3.2**: The separation of the information / data streams collection mechanisms from the search engine is a good practice for achieving resilience and meaningful results even in locations and use cases that are not sensor saturated. In such cases applications would be still able to leverage virtual sensor data feeds (such as information derived by Twitter and other on-line social networks).

5 Best Practices and Lessons Learnt on Real-time Search and Personalised Suggestions

5.1 Real-time Indexing and Retrieval

SMART tackled the major technical challenges for the real-time indexing and retrieval of social and sensor steams with the aim to serve users with the freshest results while continuously indexing data streams from SMART edge nodes. We have developed and thoroughly benchmarked a novel distributed framework, SmartReduce, which is scalable to large amount of data streams. The framework builds upon open source Storm, a reliable and transparent parallel stream processing platform, which acknowledges the wider open source strategy of the SMART framework.

5.2 Local Event Retrieval and Venue Suggestion Models

We have developed a novel event retrieval framework that is capable of identifying and ranking local events in a response to a user query (the query can be implicit or explicit). The retrieval framework has been integrated in SmartReduce and combines evidence from the content of geotagged social media and the change in the sensor observations (both physical and social) to accurately identify and rank local events.
For venue suggestions, we have also developed an anticipation component, which infers the future popularity of a venue by predicting its level of attendance using state-of-the-art time series forecasting models, relying on social data gathered from the Foursquare social network. Furthermore, we proposed a model for personalising the suggestions, by either taking into account the history of visits of the user or his/her Facebook profile.

5.3 Recommendations

Based on our experience in developing the real-time search and venue suggestion frameworks, the following recommendations for solution providers can be drawn:

- **Recommendation 5.1:** Adopting open source platforms Storm and Terrier, and making the suitable extensions to address the real-time indexing and retrieval challenges have shown the power of the open source paradigm, as well as the maturity of the used open source platforms. Since Storm and Terrier are established open source technologies with growing communities, we have indeed focused our efforts on addressing the real challenges in SmartReduce for building effective and efficient retrieval models instead of, for example, building a real-time scalable messaging system required for real-time stream processing (which is offered by Storm).
- **Recommendation 5.2**: We have thoroughly benchmarked SmartReduce for efficiency and scalability using live streams from the SMART deployed sensing infrastructure and using realistic query loads. These benchmarks show that the SMART technology is ready to be applied in real-world environments. Furthermore, it shows the importance of conducting such benchmarking procedures to ensure the efficiency

and the scalability required for real-world deployments.

- **Recommendation 5.3:** Geo-tagged social posts, such as tweets, can be used to retrieve local events around the city. However, our retrieval model's overall performance varies depending on the nature of the events. We have found that popular events obtained by crowdsourcing are easier to detect than more localized events obtained from local newspapers. . In other words, social media may not sufficiently cover high localized events that may be better identified using local sensor feeds such as crowd levels obtained from the visual processing of CCTV cameras. Indeed, we showed, with an effective modeling of background information, that such feeds would be able to detect events at the scale of a University building.
- **Recommendation 5.4:** We showed that the level of attendance of a venue could be accurately predicted using state-of-the-art time series forecasting models, and that the popularity of a venue is highly correlated with its level of attendance. However, these models are also sensitive to the sparsity of the observations, which means that, in real applications, developers need to ensure that their data streams are comprehensive, and to consider combining several sensing streams.
- **Recommendation 5.5:** For personalised venue suggestions, we have developed several models and demonstrated that the interests of a user could be inferred from either his/her Facebook's "likes", or his/her explicit preferences of venues. We have thoroughly evaluated our models, using a standard large TREC dataset and a large user study, involving 100 participants. We found out that the combination of the popularity and interestingness components are highly effective. Most interestingly, our results show that tourists generally prefer popular venues, while residents of a city prefer places that match their personal interests. Moreover, users tend to prefer personalized suggestions in the morning, while they tend to go to popular venues in the evening. Future smart urban environments should vary the level of personalisation throughout the day when recommending venues and also take into account the different personas of users (different groups of users).

6 Best Practices and Lessons Learnt on Acoustic Processing in Urban Environment

6.1 Acoustic event processing in SMART

In the SMART project we have developed the technology for acoustic event processing. The main focus of this technology is the analysis of audio that was captured in urban environment. The technology allows the classification of various common noises such as traffic and crowds. The development of this technology is based on collection of real-life audio data. This collection process is subject to the privacy restrictions that were described in section 2.

6.2 Recommendations

- **Recommendation 6.1:** Keeping the privacy and avoiding eavesdropping should be both considered from the initial stages of the audio collection planning. Those issues might cause considerable delays to the audio collection. In many cases this audio data is needed for the development process. In those cases it is best to plan for alternative data sources (e.g. commercial data) which would be used while waiting for the real data.
- **Recommendation 6.2:** Usually, the requirements from the event detect system are for "interesting" events (e.g. car crash, gun shot). However, most of the collected audio data contains only normal daily activity. The "interesting" events are usually rare and it is hard to locate them. It is best to consider this issue when planning the type of events that the system will support.

7 Best Practices and Lessons Learnt on Visual Processing

7.1 Visual processing in SMART

In the SMART project we have developed two systems, one for indoors operation performing face tracking, and another suitable for both indoors and outdoors operation performing visual scene analysis. Both systems can give an indication of how crowded the monitored space is, the way people and vehicles move and object counts through 'gates' of interest.

7.2 Recommendations

- **Recommendation 7.1:** Observing the strict privacy concerns for outdoors camera deployment in most European countries can be difficult. To get a license by the Data Protection Agency one has to plan for a time-consuming process and for a deployment where the video feed cannot be used to recognize a person, even by a human observer. This is ensured by placing the camera at a significant height, overlooking the ground.
- **Recommendation 7.2:** Again, due to the privacy concerns, the video processing algorithms need to operate on-line, since no recording of the video feed is allowed. Parallelisation of the algorithms in threads is a must to achieve processing that is fast enough.
- **Recommendation 7.3:** Face tracking in a visually cluttered environment must be based on multiple measurement cues. In SMART we have employed colour matching measurements, object detectors and foreground segmentation to build a robust system.

8 Conclusions

For more than two years SMART has been developing, deploying and validating a versatile search framework, which enables the development, deployment and operation of applications over a multitude of sensor and social networks in urban environments. The development of the SMART framework has been based on the integration of a number of components technologies spanning the areas of middleware for data collection and filtering, visual signal processing, audio signal processing, real-time information indexing and retrieval, personalized recommendations and more. There are several items that require attention regarding the deployment of each one of these components in a search application for urban environments. In particular:

- Application developers have the options of choosing between verbose high-overhead formats for data modeling which enable the capturing of rich semantics and simpler low-overhead formats with poorer semantics which however boost performance and achieve lower latency.
- In the development of SMART-like applications, application developers and integrators can leverage a wide range of data and metadata sources, including data/metadata from sensors, sensor processing algorithms, the LinkedData cloud, as well as data streams inferred from the fusion and combination of the above listed data sources.
- The design of acoustic event processing algorithms should right from the beginning consider the type of events that should be extracted, especially given the fact that in most cases end-users are interested for rare events (e.g., accidents).
- The design and implementation of video processing components should consider the ever important privacy requirements for urban deployment. These requirements impose a wide range of constraints associated with the speed of the algorithms (e.g., on-line processing is needed), as well as the deployment of cameras.
- SMART has demonstrated the power of available open source technologies (such as Terrier and Storm) in building real-time indexing and retrieval solutions for urban environments. The relevant project outcomes/developments (i.e. SmartReduce) have been benchmarked and proven to be robust and scalable.
- Personalized recommendations associated with venues (in the urban environment) could be provided based on the level of attendance of a venue. The latter level can be accurately predicted using state-of-the-art time series forecasting models, as the popularity of a venue tends to be highly correlated with its level of attendance.

Several of the above-listed technical/technological recommendations and conclusions have been driven by privacy requirements and constraints. In order to deploy and use sensors (such as cameras) in urban environments, but also in order to take advantage of social networks processing, these privacy requirements (including compliance to regulations) should be respected. The Data Protection Authorities (DPA) across EU countries are the natural contact points for discussing and consulting on privacy issues and associated requirements. Therefore deployers are advised to closely collaborate with DPAs prior to designing / planning their deployments.