

**In this Issue:**

SMART's first Open Source release now available!

SMART Workshop at Future Internet Assembly in Dublin, May 7th

SMART Technical Review

Audio Processing Advances

Video Processing Advances

Reasoning Layer

Search Layer

SMART in the Research Community

**SMART in a Nutshell**

The SMART project ([www.smartfp7.eu](http://www.smartfp7.eu)) aims to conceive and implement a scalable open source multimedia search engine that will be able to search information stemming from the physical world, Linked Data and social networks. The SMART framework will enable answering of queries based on the intelligent collection and combination of sensor generated multimedia data.

The project is co-funded by the European Commission, under the 7th Framework Program. This newsletter aims at staying in touch with relevant stakeholders and at providing updates about the project developments.

Follow us and join the SMART community!

Twitter: <https://twitter.com/smartfp7>

LinkedIn: <http://linkd.in/11tZUko>

**SMART's first Open-Source release now available!**

The SMART consortium is pleased to announce the availability of the open source SMART integrated framework! After a short registration at the following link, you will be sent full access to our release: <http://www.smartfp7.eu/content/download-survey>

The major objective of this release is to showcase the SMART technology and to attract a wide range of early adopters such as application developers, smart cities' practitioners and sensor network providers. In particular, the first open source release provides the public access to the integrated SMART framework. The release comprises of the source code and the accompanied comprehensive documentation to build, install and deploy the various components. Moreover, it offers fully functional deployments (a sandbox), hosted by the consortium in AIT and the University of Glasgow, where the integrated SMART framework can be tested. The first release is an important initial step towards building a SMART community allowing users and developers to have a feel of what the SMART framework can offer. It also gives them an opportunity to provide feedback to the consortium and the community at an early stage, which can be then used to improve the subsequent releases of SMART.

The SMART consortium has also devised a management plan for the effective monitoring and engagement to sustain and grow a SMART open source community. In particular, the consortium will continue to improve the SMART open source portal<sup>1</sup> with the required tools to effectively engage the developers and to monitor the growth of the community. In addition to providing a collaborative environment to continuously update the documentation using a wiki platform, the portal has a forum tool that allows the community to share their experiences with SMART and discuss problems they encounter and solutions they may develop when using SMART.

Subsequent versions of the SMART open source framework are expected to provide more sophisticated functionalities, while at the same time being more robust. Furthermore, later this year, application developers within the SMART consortium will integrate the present release of the SMART open source framework into more sophisticated applications associated with live news and security/surveillance use cases. The experience of developing these applications will also provide insights in fine-tuning and enhancing the open source SMART framework as part of its

<sup>1</sup> <http://opensource.smartfp7.eu>

subsequent releases.

Overall, the release of initial SMART project outcomes as open source paves the ground for continuous improvement, but also for the sustainability and the wider dissemination of the SMART research outcomes.

## SMART Workshop at Future Internet Assembly in Dublin, May 7th

Join us on May 7th in Dublin for a workshop organized by SMART and fellow FP7 projects. **“Social Networks, IoT and Cognition as enabling technologies for Smart Cities”** will explore the role of social media, IoT and cognitive technologies as essential enablers of emerging ICT-services in smart cities, while also unveiling the trend of combining them in the Cloud. We will present and discuss results, lessons learned and requirements coming out from recent and on-going EC funded projects, which intend to exploit, deploy and fully leverage those technologies in the smart cities context. The morning workshop will offer an extensive look at these domains, as well as providing networking opportunities.

Visit our website to register! ([www.smartfp7.eu](http://www.smartfp7.eu))

## SMART Technical Review

The first technical review of the SMART project was held in Brussels, December 11th, 2012. The SMART consortium members had the opportunity to present their work in terms of the SMART architecture, the open source implementation of the SMART framework, as well as the project’s dissemination and exploitation achievements. The consortium presented several demonstrations spanning individual audio-visual processing components (i.e. crowd analysis, acoustic event classification) as well as the integrated SMART search engine for environment generated content.

The initial review of the SMART project was positive and constructive. It was acknowledged that the SMART search engine and the edge server architecture are the stronger points of the project (from a technical perspective). At the same time, constructive feedback was given for improving the project’s work and final outcomes. These comments concerned the reporting of audio-visual processing advancements of the project, as well as their integration into the overall search architecture of the project. Furthermore, they underlined the need for validating the requirements addressed by SMART based on closer interaction with external stakeholders. Building on the reviewers’ comments and recommendations, the SMART consortium will continue its work in 2013 with a primary focus on producing its open source release and subsequently on using it for a wider impact.

## Audio Processing Advances

The research effort on audio processing is currently focusing on audio event classification. For this we’ve assembled an audio data

set containing about 3 hours of audio samples from the FreeSound.org web site, containing various outdoor recordings. Those recordings were manually annotated so they will be useful for training and testing the audio classifier. Our current work has focused on detecting and classifying the following audio classes: crowd noises, applause, traffic noises and music.

The audio classification is based on spectral features and uses a deep neural network (DNN) for classification. We have selected the DNN approach since it has recently been proven to be very successful in speech processing. Using this classifier and a novel training procedure, we were able to achieve an average classification rate of about 84% for the different audio classes.

Those results were verified by using two other classical classifiers: a support vector machine (SVM) and a Gaussian mixture models (GMM). We found that the DNN outperformed both of them. The classification rates for the GMM were the least successful. The SVM had classification rates closer to the DNN. We’ve also found that additional improvement can be gained by fusing the decision from the DNN and the SVM classifiers. The classification rates for the fused classifier were about 85%.

Speaker verification is another direction of research currently taking place. An improvement to the speaker verification system was achieved by using diffusion maps. This is a machine learning technique for learning manifolds in a high-dimensional space. By using this we got 20% improvement in EER compared to state-of-the-art i-vector PLDA and 40% improvement when fused with state-of-the-art i-vector PLDA.

## Video Processing Advances

The visual processing efforts have focused on face tracking and crowd analysis.



Face tracker.

Our face tracker is based on fusing three different measurement cues, colour matching, face detection and foreground information. Recently we have greatly improved our early (February 2012) face tracker from following a face robustly during 95% of the video frames and when doing so having an error of 31% of the face size, into one that always follows the face robustly with an error of 19%. This has been achieved by improving the robustness of our

measurements and the proposal distribution of the particle filter tracker. See our face tracker in action in the SMART YouTube channel<sup>2</sup>.



Crowd analysis.

Our crowd analysis system analyses crowds both statically and dynamically. The static crowd analysis system reports crowd density and primary colours in different zones of interest, using a robust adaptive foreground segmentation algorithm. Dynamic crowd analysis yields the level of compliance of the crowd motion to different directions of interest defined for each monitored zone. The latest (December 2012) system is demonstrated in our YouTube channel, as well<sup>3</sup>.

## Reasoning Layer

The reasoning layer consists of a so-called Intelligent Fusion Manager (IFM) that is meant to (i) infer high-level information from low-level information (i.e. formatted sensors data) using some patterns and (ii) to learn those patterns. We have identified the following essential requirements with regards to reasoning:

- The recognition patterns must be rule-based to ease their definition by human operators.
- The reasoning must be able to deal with uncertainty. Indeed, sensors' data may come up with some incertitude or some confidence.
- The reasoning must be able to deal with inconsistencies. For example, two sensors may produce inconsistent data.
- The reasoning must be able to deal with incomplete set of data. If a sensor or the transmission of some data is broken then the reasoning process shall be able to move on.
- The reasoning must be defeasible in the sense that new pieces of data may invalidate conclusions previously derived.

To meet these goals and requirements, we have chosen to build the IFM around Markov Logic Networks and an existing open-source implementation called Alchemy<sup>4</sup>. Finally, ethical issues are a

<sup>2</sup>Face Tracker:

<http://www.youtube.com/watch?v=MfvCh9NjldU>

<sup>3</sup>Crowd Analysis:

<http://www.youtube.com/watch?v=qG61q3TEg8Y>

<sup>4</sup><http://alchemy.cs.washington.edu/>

major concern in the Project SMART. For this reason, we are investigating another logic (a probabilistic rule-based defeasible logic) that shall ease the implementation of legal rules and privacy-abided practices.

## Search Layer

We have continued to improve our indexing and retrieval capabilities in the SMART framework, i.e. the SMART search layer. Most notably, we have delivered the first version of SmartReduce, which is the backbone of the SMART search layer. SmartReduce addresses the major technical challenges of real-time indexing and retrieval in SMART. In particular, we adapted and extended Terrier to operate with fast in-memory indices to satisfy the real-time requirement of answering user queries with the freshest results while continuously indexing data streams from the SMART edge nodes and social media. In addition, SmartReduce relies on a novel distributed framework, which makes it scalable to large amount of data streams. This is built upon open source Storm,<sup>5</sup> which is chosen over other emerging frameworks of parallel stream processing for its transparency, reliability and the growing open source community behind it. Using open source Terrier and Storm acknowledges the wider open source strategy of the SMART framework. We have also benchmarked the efficiency and the scalability of SmartReduce for indexing and answering queries using a very large number of simulated edge nodes (over 2000). The benchmarking results suggest that SmartReduce is on track to fulfil the real-time constraints that are required for a live deployment of the SMART integrated framework, which requires handling too much data in real-time from the large number of sensor streams. Specifically both the indexing latency and the query response time constraints were satisfied.

We have also made progress in developing specialised retrieval models that can identify and rank local events using evidence from multiple sensor signals and geo-located social media content. In particular, we have developed an event retrieval framework, which exploits the content of the tweets and unusual patterns in the tweeting activity in a certain area (which can be seen as a sensor signal) to identify and rank events in response to a user query. An event, in this case, is characterised by a location (in a city) and a starting time. We have performed an evaluation of this framework using a dataset of geo-located tweets from the urban area of London and local events that have occurred in London collected from local news sources and via crowd sourcing. For this evaluation, we have devised an evaluation methodology inspired by the conceptually similar problem of video segmentation task. Our empirical results show the promise and effectiveness of our framework in identifying and ranking local events. The framework can be extended to identify events from real sensor signals, e.g. the crowd/noise level, in addition to the tweeting activity signal. The sensor data that is being collected from public areas in the city of Santander will be used to validate and evaluate future versions of

<sup>5</sup><http://storm-project.net/>



this framework. For more details about the retrieval framework and the evaluation, refer to our recent OAIR 2013 paper,<sup>6</sup> available on our website.

Finally, with regards to supporting *running queries* in SMART, we have adapted and extended a traditional information filtering method to address the problem of real-time adaptive filtering of tweets. Our adaptation accounts for the main challenges that are prevalent in Twitter such as the shortness of tweets. We thoroughly evaluated our methods using the new Filtering task of the TREC 2012 Microblog track (TREC is a well-known international initiative for the evaluation of search technology).

## SMART in the Research Community

The University of Glasgow has attended two major information retrieval conferences: SIGIR and TREC. The ACM SIGIR 2012 Conference was held in Portland, USA. Iadh Ounis and Craig Macdonald have attended and presented a full paper on dynamic pruning for efficient online query scheduling. Moreover, they have presented a paper about the SMART open source framework at the Open Source Information Retrieval (OSIR) workshop collocated with SIGIR. The Text Retrieval Conference (TREC) is an annual international information retrieval evaluation forum, which has been running since 1992 by the American National Institute of Standards and Technology (NIST) to support research within the information retrieval community. The University of Glasgow has participated in various tracks of the TREC 2012 conference. Iadh Ounis has attended the conference and presented our SMART work in the real-time filtering task of the Microblog track, which has achieved one of the top performances in the track.

DEON and JURISIN are a leading international conference and workshop on legal reasoning. There, Imperial College presented new advances on a probabilistic defeasible logic and machine learning as an alternative of Markov Logic Networks to ease the integration of legal knowledge within the SMART reasoning engine. This investigation has also been published in the journal of Artificial Intelligence and Law.

Telesto was in the Mobile World Congress 2013 (Barcelona). During the matchmaking event that took place on February 26, Telesto had the opportunity to present the SMART project and meet potential cooperation partners from Research and Industry. There was significant interest and, while the project is in an early phase, the opportunity to download code and documentation was seen as very attractive.

ISCRAM is a conference regarding holistic crisis management systems. It aims at the interdisciplinary development and design of *information systems that enable better crisis response, planning*

and preparedness, mitigation and training by following integrated approaches that combine organisational, behavioural, technical, economic and environmental aspects. SMART will be presented at the ISCRAM 2013. The presentation will be about "Ethical Challenges of participatory sensing for crisis information management".

Retrieve all our publications at our website: [www.smartfp7.eu](http://www.smartfp7.eu)

## SMART Partners



**AIT**  
CENTER  
OF EXCELLENCE  
FOR RESEARCH  
AND EDUCATION

Perceptum ex Optimis



AYUNTAMIENTO DE  
**SANTANDER**



**Atos**

**IBM**

**Imperial College  
London**

**PRISA DIGITAL**

**S<sup>3</sup>log**

**Telesto  
Technologies**  
Information Communication Technologies



**University  
of Glasgow**

<sup>6</sup> M-Dyaa Albakour, Craig Macdonald, Iadh Ounis. "Identifying Local Events by Using Microblogs as Social Sensors". In proceedings of the international conference on Open research Areas in Information Retrieval (OAIR 2013)