Multimedia Search and Retrieval over Integrated Social and Sensor Networks

Paul Moore
Head of Media
AtoS Research & Innovation
Madrid, Spain
paul.moore@atosresearch.eu

John Soldatos
Associate Professor
Athens Information Technology
Peania, Greece
jsol@ait.edu.gr

***Abstract*—This paper presents work in progress associated with the development of a multimedia search engine over content and information stemming from the physical world i.e. derived through visual, acoustic and other sensors. Among the unique features of the search engine is its ability to respond to social queries, through integrating social networks with sensor networks. Motivated by this innovation, the paper presents and discusses the state-of-the-art in participatory sensing and other technologies blending social and sensor networks.**

Social Networks; Multimedia Sensor Data; Sensors; Search Engine; Multimedia Indexing; Participatory Sensing

I. Introduction

The Future Internet will include a proliferating number of internet-connected sensors, including cameras and microphone arrays [1]. Based on these sensors, emerging applications will be able to collect, filter, analyze and store large amounts of data captured from the physical world, as well as related metadata captured as part of perceptive multimedia signal processing algorithms. The ability to search this information in a scalable, effective, real-time and intelligent way can empower a wide range of added-value applications in the areas of security/surveillance, smart cities, social networking, e-science and more. The potential is partly manifested in the recent wave of participatory sensing and crowd-sourcing applications (e.g., [2], [3], [4], [5]). Nevertheless, the vast majority of crowd-sourcing and participatory sensing applications deals with non-AV (Audio-Visual) data and do not provide capabilities for searching and processing multimedia data. In addition to their inability to handle multimedia data, tools and techniques for searching sensor data (e.g., [6], [7], [8]) are still largely based on the indexing and searching of apriori defined (and usually textual) metadata. Indeed, while they exploit recent advances on sensor ontologies (e.g., [9]) in order to decouple the queries from the low level details of the underlying sensors, they cannot dynamically identify the sensors that are appropriate for answering queries according to the context of the user and the application domain.

As a result, there is a pressing need for the a next generation multimedia search engine, which will be optimized not just for textual search and text based indexing, but also for searching multimedia data (notably audio and video streams) derived from the physical world (i.e. environment generated media and content). This search engine should be able to ask queries to sensors that provide multimedia streams such as cameras and microphone arrays. In order to effectively access and use multimedia data, such a multimedia search engine should automatically match the application context with the number of sensors and sensor processing algorithms that are appropriate for answering the query at hand. To this end, cutting edge sensor processing algorithms (notably audio and video processing algorithms) can be exploited in order to allow the real-time matching of multimedia data to a given search query or application context.

In addition to integrating multimedia signal processing algorithms, multimedia search engines for environment generated content should be able to support ambient/intelligent synthesis of related content (i.e. similar to a "MapReduce" framework but customized for multimedia data). This synthesis could be based on the combination of related environment generated content and metadata, which pertain to the same context [10]. To this end, new indexing and retrieval architectures for multimedia data are required. Another challenge is the need to deal with the dynamism of the sensor web environment, where sensors, multimedia data and application contexts may change dynamically (e.g., in the scope of mobile contexts). Last (but not least) multimedia search engines for environment generated content must blend with the existing and emerging wave of social networks, in order to enable social queries over sensor networks [11].

The purpose of this paper is twofold. First, it reviews the key technological areas that should be advanced towards implemented a multimedia search engine for environment generated content. In particular, the areas of participatory sensing and the integration of social networks and sensor networks are discussed. Second, the paper aims at presenting the main characteristics of such a multimedia search engine, as planned for implementation in the scope of the EU SMART project. Along with the outline of the main characteristics of the search engine of the SMART project, the paper discusses some high-level validating use cases that could be used in order to showcase the added-value of the search engine. The structure of the paper is as follows: Section 2 following the present introduction reviews relevant developments in participatory sensing and outlines how SMART will advance participatory sensing frameworks. Likewise, Section 3 reviews the blending of social networks and sensor networks, while also outlining how such a blending could be realized in the scope of SMART. The main technical and technological characteristics of the SMART search engine are presented in Section 3. Furthermore, section 4 outlines use case for validating the concept, while section 5 concludes the paper.

II. Participatory Sensing

The SMART search framework will enable the implementation of search services over large scale community and participatory sensing infrastructures, which have recently attracted the interest of cities, communities and individuals. In particular, participatory sensing is a paradigm which uses individuals and communities to gather information about their environment using the sensing capabilities of devices such as mobile phones and car satellite navigation systems. It usually leverages the ubiquity of smart phones as sensing devices, of cloud based services for big data analysis, service and resource discovery and application delivery (i.e. app stores) while anticipating the trend towards more powerful sensing and processing capabilities of mobile devices and social networking sites. Processing can happen locally or in the cloud and devices may act singly or participate collectively in a network. Users can semantically enrich low level sensor data with high level context capturing qualitative information such as the level of rubbish on a street or their sentiments about a particular topic. Participatory sensing applications use models based on active (i.e. user driven) and passive (i.e. automated task) participation to acquire data and give feedback ranging from real time personalized recommendations to generalized trend analysis. They tend to be single focused, vertically integrated applications making design time interoperability difficult and precluding dynamic device repurposing and goal adaptation at runtime. Applications range in scale from personal sensing typically focused on data collection and analysis to give direct personal feedback (e.g. a personal health application), group sensing where information is shared among a connected group (e.g. friends, neighborhood, office etc.), to community sensing involving large numbers of users for example bird surveys. Popular services such as Google Latitude [3] and Foursquare [4] can be seen as a form of participatory sensing. They combine mobile phone based GPS with social networking and mapping services to enable users to track friends and comment on and describe features of interest. FourSquare uses a merit based system to motivate user participation.

The Personal Environmental Impact Report (PEIR) [5] integrates real time data sources with mobile device data to inform users of their environmental impact and exposure. It employs classification algorithms and contextual annotations to infer users' activity and the associated environmental impact. Map based visualizations give direct feedback information to users while integration with a social network enables peer behavior comparison. Aside from sustainability participatory sensing has been applied to domains such as health (e.g. Project HealthDesign [12], UbiFit [13]), civic engagement and citizen science (e.g. Common Sense [14], NoiseTube [15], What's Invasive) and traffic management (e.g. Mobile Millenium [16]). Privacy is key concern in participatory sensing applications. AnonySense [17] developed a framework with privacy protecting measures which enables users to opportunistically contribute data in a reliable and anonymised fashion. The CenceMe project [18] integrates with social networking and only allows sharing of data within defined user groups. PEIR has also developed functionality to enable users to introduce noise or realistic synthetic data when they do not wish to share data. Data quality is also a critical issue due to low quality and mis-calibrated sensors and the intrinsic heterogeneity of users. Factors such as user motivation, bias, activity, willingness, privacy concerns, device usage (e.g. user is speaking), device location (e.g. stationary or moving, in a pocket or exposed), transient environmental conditions etc. will have a considerable impact in the kind of inferences which can be drawn and the goals which can be achieved.

The above projects demonstrate a strong interest in participatory/community sensing services, which is empowered by the proliferation of sensor-enabled mobile devices and user generated content. These services underpin the feasibility of the SMART vision, since they provide proof-of-concept implementations for distributed sensor based search services. Furthermore, they provide ideas for exploiting user generated content (e.g., use annotations) in the SMART platform. SMART works towards extending these services emphasizing towards the following directions: (a) Exploiting multimedia audio-visual content derived from the physical world: The vast majority of the above services are based on conventional non-AV sensors (such as GPS, temperature, RFID, medical sensor). SMART introduces the perceptual components that rely on A/V processing in order to extract sophisticated context from the surrounding environment. The analysis of visual scenes and acoustic events is expected to enable more sophisticated, automated and intelligent sensor search services, due to the available of composite and richer context beyond what a device owner can contribute. In several cases, A/V processing could also obviate the need for manual annotations of the users, which can overall increase automation and reduce human intervention. Moreover, multimedia processing can leverage microphone and camera sensors, which are starting to proliferate (especially in urban environments), (b) Providing applications that span several contexts through general technology. Most of the participatory sensing applications are focused on specific domains (e.g., traffic management, trend analysis, health care). SMART proposes techniques that could cover multiple context and domains, especially in terms of routing queries and gathering results. Instead of emphasizing on a specific domain, SMART will introduce general algorithms for query routing that can take into account multiple criteria (such as location, time, environment context and more). It is envisaged that this would cover a broad spectrum of application domains, as will be also demonstrated in the scope of the SMART proof-of-concept applications.

III. INTEGRATING SOCIAL NETWORKS AND SENSOR NETWORKS

SMART works towards the combination of sensor networks information with social networks information in order to answer sensor based queries in a more social, useful and accurate way. Indeed, information from social networks can be used to enhance the end-users' context and overall understand the context of the query in a much better way. Social networks information can be used to adapt a query for environment generated context in to the end-user's daily life. The concept is quite new, yet some motivating use cases have been discussed in [19] and include (a) identifying social acquaintances in localized areas (i.e. automatically locating people from your social network residing in a specific area through GPS devices), (b) social sensing based on noise log analysis (e.g., identifying a proper place to meet individuals from your social network) and (c) Improving daily living and health for the elderly (e.g., check the status of the friends and find shopping or walking buddies to promote the mobility of elders). Other frequently discussed use case concern the ubiquitous social networking, trend, which foresees that people's updates in micro-blogging services such as Twitter, could be published automatically via their sensors.

In general there is mutual benefit (for both sensor networks and social networks) from the convergence of the two worlds. Social networks can benefit from the fact that human activity and intent can be directly derived from sensors, which obviates the needs for explicit use input. On the other hand sensor societies could start their collaboration in a social way (i.e. based on information derived by social networks). However, even though the potential of integrating social networks with sensor networks has been identified, there are still only few applications. Furthermore, there is a lack of a disciplined framework for blending sensor information with social information, including a way to deal with privacy and trust issues [20].

SMART is developing a framework for the integration of Web2.0 social networks information within searches for sensor streams and environment generated content. To this end, the project will work out semantic representations of social networks through linking social networks ontologies to the SMART information models, and accordingly sharing the semantics of the end-users social network for use within the SMART search queries. RDF modeling enables data linking and is therefore appropriate for extending SMART information structures with social information. Note that SMART includes also multimedia perceptual components that can leverage information from A/V sensors. This increases the scope of the SMART applications comparing to the state-of-the-art applications, which are mainly focused on location sensors.

Overall, SMART is working on a framework that can model information stemming for the physical world, as well as from the world of a social's network virtual communication. The framework could be accordingly used in answering SMART queries through taking into account social network information for routing queries to relevant sensors and sensor processing algorithms, gathering results, as well as ranking/grading the queries according to social criteria.

IV. SEARCH ENGINE CHARACTERISTCS

Targeting the innovations listed in the previous paragraphs, SMART will work towards a multimedia search framework with the following technical characteristics:

- **Open and Open Source**: SMART is designed as an open framework, which is extensible in terms of sensors (e.g., camera, microphone arrays, WSNs), ontologies and semantic structures (e.g., multimedia ontologies, sensor ontologies,), as well as multimedia processing components (notably video and audio processing algorithms). Furthermore, SMART will be to a large extent implemented based on open source technologies and royalty free standards. The main components of the SMART engine will be implemented in the form of open source software over the Terrier.org search engine. SMART will accordingly attempt to create an open source community for sustaining and evolving its results.
- **Multimedia**: The SMART search engine will enable query answering based on the real-time processing of multimedia data stemming from the physical environment (such as audio and video). Cutting edge audio and visual processing components will be researched and adapted, notably in the areas of acoustic event classification and visual scene analysis. These components will be used for the SMART proof-of-concept validation.
- **Participatory and Reusable**: The very same sensor and multimedia processing algorithms will be able to contribute to multiple concurrent queries of the SMART system. Participatory sensing schemes will be researched along with ways of caching data and queries, while also dealing with mobility and sharing application contexts. Furthermore, a number of Web2.0/Web3.0[21] mashups will be implemented to allow reuse of sensor queries across multiple applications and searches.
- **Smart and anticipatory**: Based on machine learning and/or rule-based mechanisms, SMART should be able to anticipate the answers to certain queries, towards proactively responding to them. This will empower a level of intelligence, beyond self-learning and ranking algorithms used by existing search engines.
- **Social**: The SMART search engine will seamless leverage information and search results from (Web2.0) social networks in order to facilitate the interception of social networks with sensor networks, towards social applications and searches of environment generated content.
- **Scalable and Dynamic**: SMART is designed to be scalable at internet scale. Hence, the project will

research a scalable architecture for collecting, filtering, processing, caching and combining sensor data in a highly heterogeneous and distributed environment. Furthermore, SMART will dynamically provide up to date information sensed by the underlying sensor networks. To this end, it will deal with the changing context of sensors (e.g., in the scope of mobility scenarios).

- **Context-aware**: SMART enables the context-aware orchestration of sensor data and metadata towards accessing data that pertain to a given context. To this end, metadata associated with time, space, location, goals, tasks and more will be used in order to trade/negotiated the contribution of a sensor to a particular query. To this end, the project researches sensor selection protocols/algorithms, along with collaborative protocols enabling the orchestration of sensors towards a joint task.

"Fig. 1" depicts a high-level overview of the SMART search engine architecture. At the lowest level of the architecture there is a cloud of sensors that provide the physical world data. The later data are processed by a number of multimedia indexing components, including visual processing components (such as video scene analysis and crowd analysis) and audio processing components (such as acoustic event classification and speech processing). Sensor information is collected and processed by appropriate edge nodes (called "sensor edge servers"), which constitute "points-of-presence" of the SMART system. Contextual information derived from multimedia processing (at the edge nodes) is indexed in appropriate multimedia knowledge bases (ontologies), which are built up in addition to sensor knowledge bases. The information held in knowledge bases is traversed in order to perform sensor selection, as well as accessing multimedia content pertaining to the end-user's query.

The upper part of the engine performs conventional processing of user queries, which includes extracting query terms ("semantics") and using them for the sensor selection and the multimedia content fetching process. The query processing process will be empowered by the SMART multimedia indexing architecture, which will extend the indexing and retrieval architecture used by the Terrier.org search engine. A detailed description of the full range of technologies comprising the SMART engine is beyond the scope of this short paper.

V. VALIDATING APPLICATIONS AND USE CASES

SMART will be validated in the areas of:

- Broadcasting and News, in order to demonstrate its ability to give fast and flexible access to factual data.
- Security and surveillance, through applications that timely access data useful for security forces and defense agencies.

The SMART use cases will be complete applications comprising multiple configurable and dynamic queries over the SMART engine. They will combine multiple search queries into composite applications. Hence, they will comprising most of the technological elements/components of the SMART search engine including audio processing components, visual processing components, sensor edge servers deployed at multiple physical points within the smart city, identification of search terms in queries, scalable distributed processing/indexing/retrieval across multiple information/sensor sources and mashups for visualization of information. As already outlined, the main applications to be developed will be live news and security/surveillance.

The "live news" use case, is motivated by the fact that timely sensor-based access to information in the urban environment can be particularly important for new agencies, which could ask the SMART engine questions regarding the occurrence and evolution of certain events, i.e. "What is happening now?", "Which places are crowed?", "What are specific trends in the city?", "Where are riots and fights happening?" and more. The answers to these queries will be provide in the form of multimedia streams mixing multimedia data acquired from the physical world (i.e. sounds/images) along with textual data stemming from sensors and metadata steams (including social networks). Using the presentation layer capabilities of the SMART framework (such as reusable Web2.0/Web3.0 mashups) news providers could build, integrate and populate web sites, wikis, blogs or news portal with news/information stemming from the underlying sensing infrastructure.

The live news use case will enable end-users to create personalized social news portals containing dynamic life information from sensors deployed within a smart city. End-users will be able to assemble a dynamic news portal, based on a set of queries to the SMART engine and associated mashup components for visualizing them on the portal. Live news applications could typically include correlated queries under a thematic umbrella. As an example consider the live news application «Crowd and Demonstrations in the City», which could lead to a personalized blog embedding the following (sample) queries to the SMART engine:

- «Which demonstrations are happening now in the city?»
- «Provide crowd events, riots and demonstrations within the city»
- «Which of my friends are close to the demonstration places?»
- «List events occurring during the demonstrations».
- «Provide information about related events in my favorite cities»
- «What is happening now in place X of the city?»
- «List similar events in similar locations in the city during the past X months»

As already outlined, the queries will be customizable and configurable, which will accordingly allow for flexibly customization news applications.

The application concerning "*Security and Surveillance in urban environments*" is motivated by a number of witnessed tremendous terroristic attacks in urban environments during the last decade (e.g., the collapse of New York's Twin Towers (2001), the bombing of packed commuter trains in Madrid (2004), as well as the London bombings (2005)). These events have led modern cities to deploy numerous sensors (notably cameras) for security purposes. Nevertheless, in such sensor saturated urban environments it has become difficult to manually observe the sensor streams. SMART can offer a viable and cost effective alternative through enabling the answering of targeted queries, based on sensors and sensor processing algorithms that fulfill certain criteria. The objective is to detect people and/or scenes that could be considered as suspicious across certain times and urban locations.

The security and surveillance use cases will leverage sensors' information and A/V processing of environment generated content streams with a view to creating wider surveillance applications for the urban environment. The applications will be built as compositions of multiple queries to the SMART search engine. At the same time, they will facilitate the issuance of alarms and the invocation of actuating services for example:

- «List suspicious sounds and the time they occur.»
- «Identify the occurrence of specific vehicle/actor movements and events»
- «List the scenes/places where a specific suspect actor is identified?»
- «Provide the locations of all the vehicles of the defense agency.»
- «Provide a list of suspicious places.»

In the scope of the above validating use cases and applications, the project will pursue the integration with real-life social networks, in order to demonstrate the merits of the SMART approach along with the described advancements over the state of the art. SMART will take advantage of the ESKUP social network, which is managed by PRISA Digital a prominent media group in Europe, which participates in the SMART consortium. The ESKUP platform includes information and users' profiles that can be directly correlated to information stemming from sensors deployed in Spanish smart cities. Liaisons with some other big European social networks (e.g., those listed in http://www.socialnetworksgroup.eu/) will be also pursued. Note however that ESKUP presents distinct advantages for the SMART validation, given that: (a) It is thematically pertinent to the events and contexts captured from the SMART sensors, (b) It is owned and controlled by the SMART consortium partner PRISA Digital.

VI. Conclusions

This paper has presented work in progress, which is planned to be carried out in the scope of the FP7 EU co-funded project SMART. Given that the work is in its infancy, the paper has provided the main concepts and ideas underpinning SMART, along with a brief description of some of the background work and technologies. The work is focused on the development of a multimedia search framework facilitating the development of search applications that access, process and visualize information stemming from the physical world. Among the main characteristics of the framework is its ability to integrate information stemming from social networks with a view to endowing query results with a social dimension. This feature is in-line with emerging applications that blend social networks with sensor networks, as well as with the wave of participatory sensing applications. SMART put emphasis on processing and indexing multimedia information (i.e. stemming from audio-visual) data, which is a key distinguishing characteristic comparing to conventional participatory sensing applications. To this end, SMART is working on a novel multimedia indexing architecture, along with leading edge components for audio and visual processing of physical world information (such as algorithms for visual scene analysis, crowd analysis and acoustic event detection and classification). In this paper we have focused on the social and participatory dimension of the SMART project, rather than providing details of the work towards multimedia processing and indexing. We expect that the next three years will lead to the development of an open search framework for multimedia data stemming from physical world sensors. This framework would enable the development of a new breed of social multimedia search applications, which are not possible nowadays.

Acknowledgment

This work is part of the EU SMART project (proposal number FP7-287583) currently under negotiation with the EC (European Commission) in the scope of the FP7 ICT framework. The authors acknowledge valuable help and contributions from all partners of the SMART project.

References

[1] Harald Sundmaeker, Patrick Guillemin, Peter Friess, Sylvie Woelfflé (eds), "Vision and Challenges for Realising the Internet of Things", March 2010, ISBN 978-92-79-15088-3, doi:10.2759/26127, © European Union, 2010.

[2] Sasank Reddy, Deborah Estrin, Mani Srivastava. Recruitment Framework for Participatory Sensing Data Collections. In Proceedings of the 8th International Conference on Pervasive Computing (May 2010), pp. 138-155.

[3] Page, X. and A. Kobsa (2010): Navigating the Social Terrain with Google Latitude. iConference 2010, Urbana-Champaign, IL, p.174-178.

[4] Marshall Kirkpatrick "The Era of Location-as-Platform Has Arrived", ReadWriteWeb, January 25, 2010

[5] Min Mun, Sasank Reddy , Katie Shilton, et. al., "PEIR, the personal environmental impact report, as a platform for participatory sensing systems research", in the Proceedings of the 7th international conference on Mobile systems, applications, and services, MobiSys '09.

[6] Dominique Guinard, Vlad Trifa: Towards the Web of Things: Web Mashups for Embedded Devices, WWW2009, 2009, Madrid, Spain.

[7] MHildebrand, J van Ossenbruggen and L Hardman, An Analysis of Search-based User Interaction on the Semantic Web, Tech. rep., Centrum voor Wiskunde en Informatica (NL), 2007.

[8] J. Camp, J. Robinson, C. Steger, and E. Knightly. Measurement driven deployment of a two-tier urban mesh access network. In Proc. ACM MobiSys 2006, Uppsala, Sweden, June 2006.

[9] Daniel 0 'Byrne, Rob Brennan, Dec/an 0 'Sullivan: Implementing the Draft W3C Semantic Sensor Network Ontology, In 8th IEEE International Conference on Pervasive Computing and Communications Workshops (PERCOM Workshops), 2010.

[10] Mark Ingebretsen, "Making Sense of Sensor Data," IEEE Intelligent Systems, vol. 25, no. 3, pp. 3-5, May/June 2010, doi:10.1109/MIS.2010.86.

[11] John G. Breslin, Stefan Decker. The Future of Social Networks on the Internet: The Need for Semantic. IEEE Internet Computing, Volume 11, Issue 6 (Nov/Dec 2007), pp. 86-90.

[12] Patricia Flatley Brennan “Project HealthDesign: Rethinking the power and potential of personal health records”, Journal of Biomedical Informatics archive, Volume 43 Issue 5, October, 2010.

[13] Consolvo, S., McDonald, D., Toscos, T., Chen, M.Y., Froehlich, J., Harrison, B., Klasnja, P., LaMarca, A., LeGrand, L., Libby, R., Smith, I., & Landay, J.A. Activity Sensing in the Wild: A Field Trial of UbiFit Garden. In Proc. CHI 2008, ACM (2008) 1797-1806.

[14] Catherine Havasi, James Pustejovsky, Robert Speer, Henry Lieberman, Digital Intuition: Applying Common Sense Using Dimensionality Reduction, IEEE Intelligent Systems 24(4), pp. 24-35, July 2009.

[15] Nicolas Maisonneuve, Matthias Stevens, Maria E. Niessen, Luc Steels, “NoiseTube: Measuring and mapping noise pollution with mobile phones”, Information Technologies in Environmental Engineering (ITEE 2009), Proceedings of the 4th International ICSC Symposium, Thessaloniki, Greece, May 28-29, 2009.

[16] R. Herring, A. Hofleitner, S. Amin, T. Nasr, A. Khalek, P. Abbeel, A. Bayen, “Using Mobile Phones to Forecast Arterial Traffic Through Statistical Learning”, 89th Transportation Research Board Annual Meeting, Washington D.C., January 10-14, 2010.

[17] Cory Cornelius, Apu Kapadia, David Kotz, Dan Peebles, Minho Shin, and Nikos Triandopoulos. Anonysense: privacy-aware people-centric sensing. In Proceeding of the 6th international conference on Mobile systems, applications, and services, MobiSys '08, pages 211-224, New York, NY, USA, 2008. ACM.

[18] Emiliano Miluzzo, Nicholas D. Lane, Shane B. Eisenman, Andrew T. Campbell,“CenceMe: injecting sensing presence into social networking applications”, in the Proceedings of the 2nd European conference on Smart sensing and context EuroSSC'07.

[19] John G. Breslin, Stefan Decker,, Manfred Hauswirth, et. al., “Integrating Social Networks and Sensor Networks, W3C Workshop on the Future of Social Networking, 15-16 January 2009, Barcelona.

[20] A. Hogan, A. Harth, A Polleres. SAOR: Authoritative Reasoning for the Web. In Proceedings of the 3rd Asian Semantic Web Conference (ASWC 2008), volume 5367 of Lecture Notes in Computer Science, pages 76-90, 2008.

[21] Jim Hendler, “Web 3.0 Emerging”, in IEEE Computer, Vol. 42, No. 1. (January 2009), pp. 111-113.

[22] H. Goto, Y. Hasegawa, and M. Tanaka, “Efficient Scheduling Focusing on the Duality of MPL Representatives,” Proc. IEEE Symp. Computational Intelligence in Scheduling (SCIS 07), IEEE Press, Dec. 2007, pp. 57-64, doi:10.1109/SCIS.2007.357670.

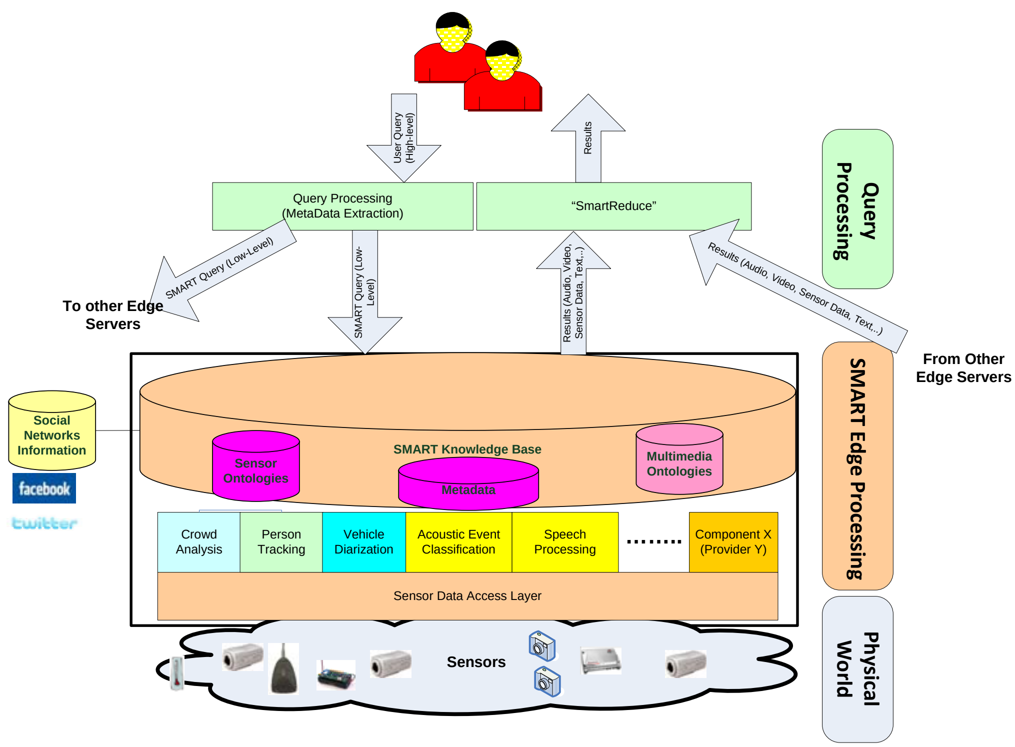

Figure 1. Anatomy of the SMART Multimedia Search Engine