Urban context detection and context-aware recommendation via networks of *Humans as Sensors*

Sergio Alvarez-Napagao, Arturo Tejeda-Gómez, Luis Oliva-Felipe, Dario Garcia-Gasulla, Victor Codina, Javier Vázquez-Salceda

Universitat Politècnica de Catalunya, Barcelona, Spain
{salvarez,jatejeda,loliva,dariog,vcodina,jvazquez}@lsi.upc.edu

Abstract. The wide adoption of smart mobile devices makes the concept of *human as a sensor* possible, opening the door to new ways of solving recurrent problems that occur in everyday life by taking advantage of the information these devices can produce. In the case of this paper, we present part of the work done in the EU project SUPERHUB and introduce how geolocated positioning coming from such devices can be used to infer the current context of the city, *e.g.*, disruptive events, and how this information can be used to provide services to the end-users.

Keywords: social networks, smart mobile devices, human as a sensor, recommender systems

1 Introduction

Mobility is one of the main challenges for urban planners in cities. Even with the constant technological progress, it is still difficult for policy makers and transport operators to 1) know the state of the city in (near) real-time, and 2) achieve proximity with the end-user of such city services, especially with regards to communicating with the citizen and receiving proper feedback.

There is a relatively recent technological advance that enables an opportunity to partially tackle these issues: ubiquous computational resources. For instance, thanks to smartphones, users that move in a city can potentially generate automatic data that may be hard to obtain otherwise: location, movement flow, average trip times, and so on. Moreover, transport network problems and incidents that affect mobility services are often documented by someone somewhere in the Internet at the same time or even before, than they appear in official sources or in the news media. This phenomenon has been referred to as *humans as sensors* [14]. Sensing through mobile humans potentially provides sensor coverage where events are taking place. An additional benefit is that human expertise can be used to operate such sensors to raise the quality of measurements, through *e.g.*, a more intelligent decision making, such as setting up a camera in an optimal way in poor lighting conditions; or providing exploitable additional metadata, as in collaborative tagging processes such as hashtagging.

In this paper, we show a system that is able to mine such data in order to:

1. improve knowledge obtained from other data generation approaches, such as GPS pattern analysis,
2. detect unexpected situations in the city that may affect large groups of people at a certain location, *e.g.*, public demonstrations or celebrations, sudden traffic jams caused by accidents, and
3. enable services to users that exploit such generated knowledge, providing novel kinds of real-time information and recommendation.

The paper presents, due to space constraints, just a general overview of the problems we tackle, the preliminar results of the parts already implemented, and the future work. For deeper reports on the technical details, please refer to the related deliverables[1] and to [8].

This paper is structured as follows: in §2 we introduce SUPERHUB, an urban mobility-related EU project; §3 contains an explanation of the extent of the contextual detection, focusing on social network data; §4 explains how the contextual information generated can be used to provide services for the end-users; and finally §5 presents related work and wraps up the paper with conclusions.

2 The case of SUPERHUB

SUPERHUB [2] is a project co-funded by the European Commission. Its main goal is to provide an open platform capable of considering in real time various mobility offers, in order to provide a set of mobility services able to address users' needs. At the same time the project intends to promote user participation and environmental friendly and energy-efficient behaviours.

To achieve these objectives SUPERHUB is developing, but not limited to (see Figure 1): 1) novel methods and tools for real-time reasoning on large data streams coming from heterogeneous sources; 2) new algorithms and protocols for inferring traffic conditions from mobile users, by coupling data from mobile phone networks with information coming from both GPS data and social network streams; and 3) a journey planner with the goal of best fulfilling user mobility needs and preferences while minimizing negative environmental impact.

The project builds on the notion that citizens are not just mere users of mobility services, but represent an active component and a resource for policy-makers willing to improve sustainable mobility in smart cities. Existing journey planners only provide a few options to let users customize, to some extent, how the journey should look like. The reality, however, is more nuanced – different users might prefer different routes which, in addition, depend on the users context (*e.g.*, a shopping trip, travelling with small children or going back home) as well as on the environmental context: weather, traffic, crowdedness, events, *etc.*

At a technical level, this is reflected in a front-end component of the system that automatically builds, maintains and adapts users profile over time, which includes a detailed description of users preferences in terms of mobility options.

[1] http://www.superhub-project.eu/downloads/viewcategory/6-approved-deliverables.html

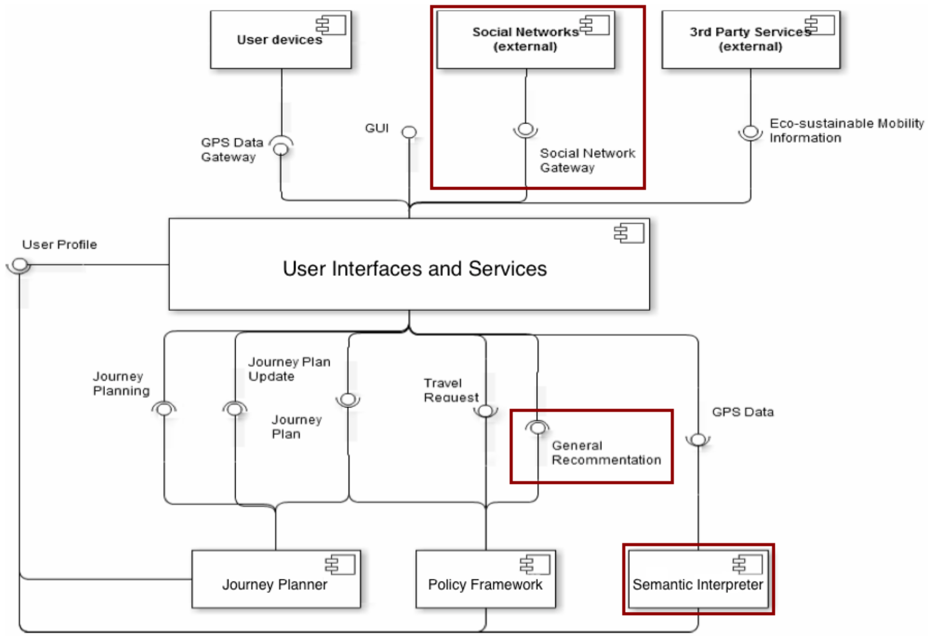

Fig. 1: *Persuasive mobility Services and Interfaces* view of the SUPERHUB architecture [4]. In boxes, the components this paper focuses on.

The profile is used to customize and rank mobility offers, in such a way to promote solutions as close as possible to users expectations, in terms of their needs and goals, while fostering the adoption of environmentally-friendly offers.

By using 1) the contextual information obtained by a wide range of heterogeneous sources and 2) such user profile, SUPERHUB is able to rank mobility offers and provide personalized recommendations.

3 Real-time detection of urban context

In a real-world setting, the SUPERHUB journey planner has to be ready to receive high amounts of journey plan requests and deliver multimodal recommendations that best fit a wide range of criteria, including user preferences. However, the evaluation of such criteria is continuously dependent on factors that occur in the external world, what we call the context of the request. And the context, given the same external conditions, is always city-dependent.

For example, if a request for a journey plan is made in January while it is snowing, in the case of Helsinki it is considered a normal situation and public transport should not be affected by such weather conditions. However, if the user is in Barcelona it is considered an extraordinary case that has numerous, unpredictable consequences on the transport networks. Therefore, weather sensor data may be useful to detect a context, but is not sufficient by itself.

3.1 From heterogeneous sensor data to real-time knowledge

This process is being carried in SUPERHUB by the Semantic Interpreter, taking advantage of semantic interpretation techniques in order to infer knowledge from both situational and historical data. Such knowledge can be applied by other components for diverse purposes, such as generating more fine-grained user models, or being able to understand normality with respect to policy fulfillment and thus derive and predict unexpected situations.

The Semantic Interpreter is a component that provides knowledge in the form of RDF triples inferred from sensor data that is of a higher level of abstraction than what is usually obtained with other techniques, acting as a central point for data homogeneisation. Via the Semantic Interpreter, raw data is filtered, normalised and interpreted into high-level concepts. Such concepts can be merged and analysed to generate derivative concepts that are not explicit in the sensor data but implicit in the aggregation of large instances of it. The analysis relies in applying semantic inference via Pellet, fed by expert knowledge and constitutive information about the city in SWRL, applied to statistical aggregations.

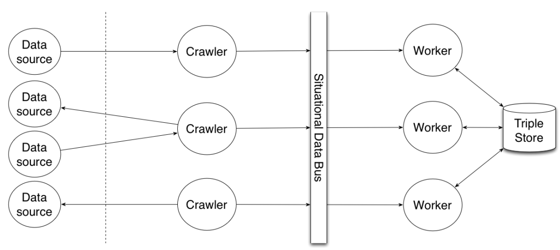

Fig. 2: Sketch of the internal components of the Semantic Interpreter.

The SUPERHUB Semantic Interpreter has been implemented in Clojure [10] (a dialect of Lisp) and has been collecting and processing data on top of the JVM since July 2013. The raw data, in a normalised form (called *Situational Data*) is stored in MongoDB as JSON documents, while the RDF produced representing high-level knowledge, *e.g.*, disruptive events, are stored in a Neo4j graph database. Internally, the implementation is based on independent, autonomous agents that communicate with each other exclusively by asynchronous messages via *Clojure agents*. The agents have the capability of proactively assign themselves a particular role (see Figure 2):

- *Crawler* agents assign themselves a target API or web service and manage the reception of data from them, coordinating between them which endpoints from a common pool should be queried or listened to next, and
- *Worker* agents schedule periodical aggregation processes. Aggregation processes can be hot-plugged and removed from the Semantic Interpreter at

runtime via plug-ins, and can include but are not limited to: crowdedness by area and time interval, crowdedness by Point of Interest and time interval, user trajectories by time interval, disruptive events detection, *etc*.

In a data-sensible application such as the Semantic Interpreter, reliability is a crucial feature. In our system agents are fail-safe in the sense that if a process fails, another agent is taken from a pool to automatically select one or more roles and fulfill them. Scalability is handled by the Semantic Interpreter by not allowing more agents than $n-1$, where n is the number of cores of the host. The system handles tasks in an abstract way, while the execution part is held autonomously by agents that can plan individually and coordinate between themselves. Therefore, the system is decentralised and scheduled tasks are always picked up, and a result of that failures in the Semantic Interpreter have been very sporadic (two shortages in 7 months).

An instance of the Semantic Interpreter can be parametrised by setting up the following values in a configuration file: latitude and longitude of the central coordinate, radius of the metropolitan area of the city, counts-as rules (city-specific interpretation rules in RDF), social network API keys, the credentials to MongoDB and Neo4j, and the periodicity of the aggregation processes. This means that, with a small setup, the instance can be applied to any city.

The RDF information generated by the Semantic Interpreter through each of the aggregation processes carried out by agents is visually represented by the SUPERHUB Situational Data Visualiser. In Figure 3 there is an example of such information: raw data from social networks (dots), density by area (rectangles), and user trajectories (arrows) corresponding to a span of 15 minutes.

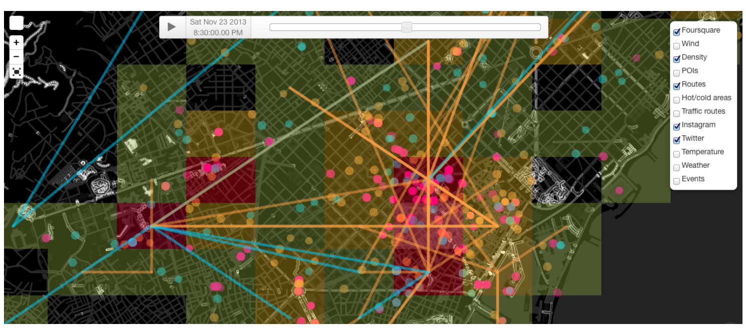

Fig. 3: Screenshot of the SUPERHUB Situational Data Visualiser, showing raw data, density by areas and user trajectories.

3.2 Inferring city context from social networks

In SUPERHUB, many types of sensors and services are used to get a picture of what is going on in the city: GPS real-time data, mobility flow patterns, official alerts, social network information propagation and weather forecasts among others. By simply combining incoming loads of such data, it is not possible to make pertinent complex decisions: we need to infer what the big picture is in order to interpret the current context and its implications. After that we will be able to detect relevant situations.

Basically, a normal situation in the city can be modelled under a spatio-temporal context which can be defined as any information that characterises a situation. Disruptive events and incidents in the mobility city field are abnormal situations that have a significant impact on the city mobility, for instance: traffic jams or metro service delays. In order to maintain a reliable mobility city behaviour, it is necessary to detect and predict these kind of events.

In the case of Twitter, tweets are aggregated and analysed in time windows in order to find located trends [15]. In information diffusion, reliability of data obtained is a big concern, and our solution is to filter bots and spammers by analysing the topology of the individual social networks of the users [1], and by calculating their influence on other users [9]. The Foursquare API adds to the reliability score of the aggregated data by contrasting the detected trends against area-based collections of check-ins in the same time windows.

Currently, the Semantic Interpreter is able to detect disruptive events based on abnormal social network activity. By grounding the model of events spatio-temporally we build a representation of each event which allows temporal and spatial reasoning. Our final event representation, further discussed in [8], allows us to know the relevance of the event (how much it deviates from expected behavior), the impact of the event (what is its effect on nearby areas and mobility services) and how the event forms and disperses (when did the event actually begin, when it reached full certainty, when did it began to fade away and for how long) among others. All those features together with the results of a collaborative tagging process (see §3.3) empowers interesting features, supported by promising results. As an example of those consider the most relevant events detected by the system (those with a higher deviation from the expected behavior). The top 40 events are shown in Table 1.

Another example of potential application of this methodology is that of *event crowd estimation*. Since Table 1 indicates that events detected with more certainty are those with more popular assistance, we considered the possibility of estimating the actual assitance to an event through the data collected by our system. We obtained the estimated assistance to the top 100 events detected from official sources when available, and estimated the rest from venue capacity. The result is a dispersion chart showing the relationship between the average activity captured by our system and the actual attendance of the event, along with a computed linear regression (see Figure 4). The Pearson correlation coefficient is approximately 0.82, which is a relevant result considering the dispersion of the data collected. This experiment suggests that we can automatically estimate the number of people at an event through its representation in our model.

1	FCB vs Madrid	11	3 nearby concerts	21	Depeche Mode concert	31	Barcelona Fashion Week
2	FCB vs Elche	12	Airport	22	New Year @ Park Guell	32	FOALS concert
3	FCB vs Malaga	13	Michael Buble concert	23	Daughter concert	33	Christmas shopping
4	FCB vs RCDE	14	Arctic Monkeys concert	24	RCDE vs Madrid	34	Biffy Clyro concert
5	FCB vs Milan	15	New Year @ Park Guell	25	Airport	35	FCB vs Real Sociedad
6	FCB vs Granada	16	Airport	26	Parc Guell visit	36	Mishima concert
7	FCB vs Valencia	17	Bruno Mars concert	27	FCB vs Getafe	37	Airport
8	FCB vs Villareal	18	FCB vs Efes (Basketball)	28	New Year @ Sagrada Familia	38	RCDE vs Madrid
9	FCB vs Celtic	19	FCB vs Real Sociedad	29	New Year @ Placa Espanya	39	FCB vs Levante
10	FCB vs Cartagena	20	Airport	30	Camp Nou visit	40	Lori Meyers concert

Table 1: Top 40 events by impact and their description. FCB stands for FC Barcelona (football, basketball) and RCDE stands for RCD Espanyol (football).

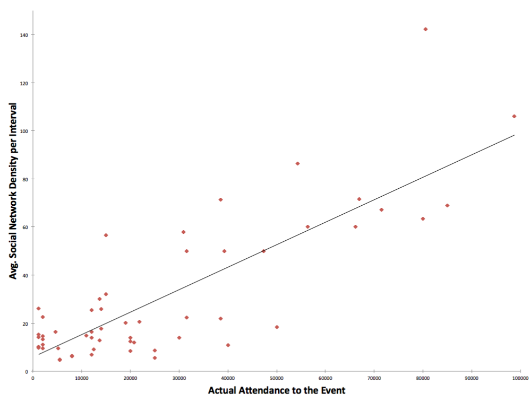

Fig. 4: Captured data vs. actual attendance for the Top 100 events detected.

3.3 User input as a basis for interpretation

When disruptive events happen – *e.g.*, accident, traffic jam, sports events, public demonstrations, flashmobs –, in most cases it will be captured by the SUPERHUB system thanks to the combination of heterogeneous data sources. However, in many cases there is no further available information that can explain exactly what happened than a geolocated position. Social networks and user inputs can be used as part of the disruptive event detection process.

We take advantage of stakeholders and users as experts that annotate and learn about detected events to identify and predict them. Users can inform and provide context about events in real-time, as well as provide tagging information of unclassified past events. In such cases, the identification of recurrent events allow the classification of what is normal and what is not. Currently, user feedback is retrieved by providing a collaborative tagging web application (see Figure 5). In this applications, users are presented random disrupted events (or they pick a specific one from the Situational Data Visualiser), and by using an autocomplete control they assign tags that are linked to RDF concepts. These concepts currently include DBPedia, WordNet RDF, and an *ad hoc* ontology of Points of Interest based on the Foursquare API.

Fig. 5: Screenshot of the app for collaborative tagging of disruptive events.

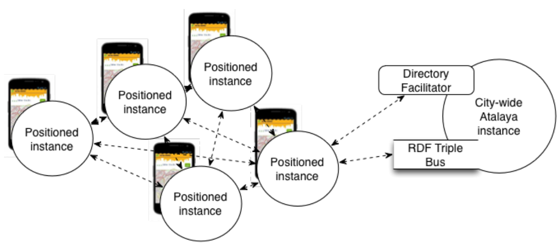

Fig. 6: Schema of a distributed system of Semantic Interpreters.

Currently, the SUPERHUB Android mobile app includes a lightweight interface that communicates with the journey planner API. One of the features that this app provides is the possibility of reporting a disruptive event, by choosing options specified in a taxonomy. The mobile app also sends, if the user explicitly authorises so, GPS traces with the positioning of the user accross time. As future work, and taking advantage of the fact that the Semantic Interpreter is easily configurable and automatically scalable, we plan to embed a small and lightweight version of the Semantic Interpreter instance, working transparently as a low-priority background process. This instance can be used to obtain more detailed data coming from the social networks by setting up a smaller radius as a parameter[2], and to perform basic aggregation processes.

Figure 6 shows an schema of this proposal: the instance responsible for the whole city is in charge of maintaining in real-time a list of all mobile app-embedded Semantic Interpreter and exposing an API for its Triple Store. Mobile-embedded instances are connected between themselves, understanding the position of the nearest ones and thus dynamically deciding whether to expand or shrink their radius of data gathering.

4 Context-aware user-centric services

The use of personalization in the field of journey planning is mainly focused on tourism [13]. Although there are other domains where routes are also recommended, *e.g.*, for sporting/leisure purposes [5,11]. For instance, [5] produces journey hiking walkways according to a set of milestones or personalize journey plans by choosing from pre-defined routes which better fit users preferences. Afterwards, the selected routes are enhanced by means of adding Points Of Interest (POI) -relevant geographical features that may be relevant to the user.

In general, these approaches can be considered as aimed at providing routes to users on closed domains, that is: 1) the set of potential journey plans is already defined, and 2) the journey plans are built based on a recommended set of POIs (*e.g.*, touristic routes). Other approaches try to overcome this closed domain constraint. While some solutions successfully generate this kind of routes, they do not incorporate users preferences or just a small set of generic preferences are taken into account (*e.g.*, departure time or cost) [12].

SUPERHUB aims to generate journey plans from a more open perspective. Users select a destination and plans are designed accordingly, using different modes of transport, including car-pooling, and considering users preferences and contextual information (*e.g.*, weather and current users situation). Consequently, the routes are not enclosed to contain certain POIs and the space search is broader and more complex. The objective of our recommender is two-fold: 1) recommending journey plans, and 2) providing opportunistic recommendations in real-time. A good example of both functionalities can be found in [5]. Our recommender differs from that approach in that routes are not taken from a pre-defined set, but are dynamically generated according to contextual situation in the city. Thus, [5] adds a set of POIs to enhance the routes which is similar to our approach for opportunistic recommendation, although it does not take into consideration contextual information.

As discussed previously in §2, the SUPERHUB recommender makes use of contextual data to perform context-aware recommendations. When context mat-

[2] This is due to the fact that the public APIs of the social networks have request limitations per each API key. Users of the mobile app have the option to provide connection to social networks, and this connection can be used to obtain more fine-grained data, virtually expanding the city-wide Semantic Interpreter capabilities.

ters, as is the case of the mobility domain – *e.g.*, when it is raining, people generally preferto avoid walking or riding a bike –, it is sensible to use a user model that has been learnt with feedback acquired in the same context as the target user, as only that feedback is relevant for the prediction.

Given the different nature of both functionalities - predicting user satisfaction towards a suggested journey plan and suggestiong appealing POIs along a route - our recommender system implements two different recommender algorithms to provide them.

In one hand, predicting user score towards journey plans; ideally, there is a large amount of possible Journey Plans or Legs. This implies that any feedback retrieved from users may be too sparse, which makes difficult to find similar users to proceed with a collaborative-filtering approach (*e.g.*, only users that perform the same routes would be found as similar users). Our solution adopts a content-based approach, which is consistent with current state of the art on mobility recommenders. By means of extracting features that define a Journey Plan or Leg, it is possible to learn the preference towards each feature from users feedback given to previous journey plans. For instance, a user may provide positive feedback towards journey plans that require little walking and mainly use the bus at night, and negative feedback to journey plans that require walking or riding a bike also when it is dark. Thus, we can predict better score to journey plans that include taking the bus when night rather than walking. It is worth noticing that user may have explicited the preference towards riding a bike or a bus, but the user model has learnt the preference about the same modes of transport under different contexts. Thus, the training process incorporates users preferences as a starting point and process feedback from previous journey plans done to refine those preferences.

In the other hand, suggesting appealing POIs along a route, our recommender filters a set of POIs by selecting those that are aligned to the user and the context. Given that the domain of POIs is not as large as journey plans we adopted a collaborative-filtering approach, which has proven to be effective and requires less information about items (POIs in this case). Nevertheless, collaborative-filtering can suffer from data sparsity, thus we chose a Matrix Factorization algorithm to find latent factors and reduce this issue. For both algorithms, their learning process and resulting user models share the same input information:

– The situational data, which includes all the wide range of contextual information occurring in the city (*e.g.*, weather, pollution, crowdedness).
– The history of previous recommendations, that is journey plans and POIs suggested, along the user's context (*e.g.*, purpose and companionship).
– The feedback given by the user about the suggested items (plans or POIs).
– The User profile, to retrieve user's preferences for bootstrapping purposes.

For each user, different user models are learnt according to the set of contexts that are more significant or allow distinguishing better what the user likes most when travelling. Then, when new recommnendations are needed, our system picks the user model that better fits the current context and uses it to either predict user's preference towards a journey plan or select the most appealing POIs that s/he may like to know while travelling.

The SUPERHUB recommender, by following approaches for context-aware recommendation [3], allows improving results and dynamically adapting to the real situation in the city and to users' needs, goals and current context (*e.g.*, in a rush, going to work, with children). In order to do so, it intensively mines the contextual information generated by the Semantic Interpreter to pre-filter user models and to learn and exploit contextual user models. This means that the training process of each recommender computes a user model for each kind of context (*e.g.*, raining, raining and night, sunny). These models are learnt offline, since computing all of them require a significant amount of time. Using this approach, having all the user models of a given user pre-calculated allows improving scalability of the recommending process.

5 Conclusions

The approach proposed in this paper is to use a combination of sentiment analysis, semantic inference, information diffusion and big data techniques on top of interactions of users in certain social networks, combined with city-specific knowledge, city events calendars and geospatial data. Data is analysed in two stages, on-line and off-line, in order to learn behavioural patterns and capture those deviations that may reflect events that affect urban mobility. The Semantic Interpreter is currently running and producing results for Barcelona, Milan and Helsinki, the three trials cities for the SUPERHUB project evaluation. By using an execution mechanism inspired by agent paradigms we make sure that the system is self-dependant and reliable enough to automatically build crowdedness models with data as volatile as social network inputs.

Depending on the periodicity of the aggregation set in the configuration of the Semantic Interpreter instance, the delay between the start of the event in the city and the instant at which it is detected by the SUPERHUB platform may vary, but we have empirically proven it can be as low as 5 minutes.

Additionally, we propose to take advantage of the mobile app by enabling it with a local semantic interpretation process providing an immediate associated functionality: enhancing the reporting of disruptive events by proactively suggesting events happening nearby the user and presenting the option to confirm or deny them, or to tag them. In this way, users would be generators not only of geolocated data, but also of first-person knowledge of their surrounding context.

Similar proposals have recently appeared that cover some of the aspects of such a system, either from the individual sensoring perspective, closer to multi-agent based systems [6] or from big data techniques applied to social network streams by combining statistical analysis with semantic interpretation [7]. However, such proposals work at a granularity level that is either too high or too low: in the former approach there is, in principle, no centralised mechanism to maintain global aggregations; in the latter, systems do not take advantage of the end-user terminals and their potential dispersion in time and space. In this paper we present a system that can work at both levels, leveraging global aggregations with local awareness to have the best possible picture of the context of the city.

Finally, the SUPERHUB recommender, by having a strong focus on the use of semantically-enhanced contextual information, provides novelty to the state of the art beyond the use of recommenders in the mobility domain.

References

1. S. Boccaletti, V. Latora, Y. Moreno, M. Chavez, and D. HWANG. Complex networks: Structure and dynamics. *Physics Reports*, 424(4-5):175–308, Feb. 2006.
2. I. Carreras, S. Gabrielli, D. Miorandi, A. Tamilin, F. Cartolano, M. Jakob, and S. Marzorati. SUPERHUB: a user-centric perspective on sustainable urban mobility. In *Sense Transport '12: Proc. of the 6th ACM workshop on Next generation mobile computing for dynamic personalised travel planning*. ACM, June 2012.
3. V. Codina, F. Ricci, and L. Ceccaroni. Local Context Modeling with Semantic Pre-filtering. In *Proceedings of the 7th ACM Conference on Recommender Systems*, pages 363–366, New York, NY, USA, 2013. ACM.
4. S. Cretti and F. Facca. D2.1 FP7-ICT-2011-7 SUPERHUB - Report on the architecture definition Open Source strategy and adoption pattern. Technical report, Mar. 2012.
5. C. Diaspero, A. Heinisch, and A. Petrova. A Mobile Recommender System for Hiking Walkways. 2011.
6. C. Ellul, S. Gupta, M. M. Haklay, and K. Bryson. A Platform for Location Based App Development for Citizen Science and Community Mapping. In *Progress in Location-Based Services*, pages 71–90. Springer Berlin Heidelberg, Berlin, Heidelberg, Jan. 2013.
7. L. Gabrielli, S. Rinzivillo, F. Ronzano, and D. Villatoro. From Tweets to Semantic Trajectories: Mining Anomalous Urban Mobility Patterns. In *Citizen in Sensor Networks*, pages 26–35. Springer International Publishing, Cham, Dec. 2013.
8. D. Garcia-Gasulla, A. Tejeda-Gómez, S. Alvarez-Napagao, L. Oliva-Felipe, and J. Vázquez-Salceda. Detection of events through collaborative social network data. *Proceedings of the 6th International Workshop on Emergent Intelligence on Networked Agents (WEIN'14)*, May 2014.
9. J. T. Gomez, M. S. Marrè, and J. P. Serra. tweetStimuli: discovering social structure of influence. 2012.
10. R. Hickey. The Clojure programming language. In *DLS '08: Proceedings of the 2008 symposium on Dynamic languages*. ACM, July 2008.
11. S. Knoch, A. Chapko, A. Emrich, D. Werth, and P. Loos. A Context-Aware Running Route Recommender Learning from User Histories Using Artificial Neural Networks. *Database and Expert Systems Applications (DEXA), 2012 23rd International Workshop on*, pages 106–110, 2012.
12. L. McGinty and B. Smyth. Personalised Route Planning: A Case-Based Approach. In *Advances in Case-Based Reasoning*, pages 431–443. Springer Berlin Heidelberg, Berlin, Heidelberg, Jan. 2003.
13. F. Ricci. Travel recommender systems. *IEEE Intelligent Systems*, 2002.
14. M. Srivastava, T. Abdelzaher, and B. Szymanski. Human-centric sensing. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 370(1958):176–197, Nov. 2011.
15. J. Weng and B. S. Lee. Event Detection in Twitter. *ICWSM*, 2011.