Semantic management of moving objects: A vision towards smart mobility

Sergio Ilarri [a,*], Dragan Stojanovic [b], Cyril Ray [c]

[a] Department of Computer Science and Systems Engineering, University of Zaragoza, Spain
[b] Faculty of Electronic Engineering, University of Nis, Serbia
[c] Institut de Recherche de l'École Navale (IRENav), Brest, France

ABSTRACT

This position paper presents our vision for the semantic management of moving objects. We argue that exploiting semantic techniques in mobility data management can bring valuable benefits to many domains characterized by the mobility of users and moving objects in general, such as traffic management, urban dynamics analysis, ambient assisted living, emergency management, m-health, etc. We present the state-of-the-art in the domain of management of semantic locations and trajectories, and outline research challenges that need to be investigated to enable a full-fledged and intelligent semantic management of moving objects and location-based services that support smarter mobility. We propose a distributed framework for the semantic enrichment and management of mobility data and analyze the potential deployment and exploitation of such a framework.

1. Introduction

With rapid advances in sensor technologies and wireless communications, a plenty of positioning methods and systems have been developed, that determine the locations of moving objects with different accuracy. A sequence of locations, represented as timestamped points in appropriate geometric/geographic space, makes a trajectory of a moving object in its raw form. A lot of research has been first performed on the modeling, representation, processing, querying, analysis and mining (Li et al., 2011; Wu, Lei, Li, & Han, 2014) of such raw moving object trajectories by the *moving object database* (*MOD*) community (Wolfson, Xu, Chamberlain, & Jiang, 1998; Wolfson & Mena, 2004). However, *smart mobility* services obviously need to go beyond this straightforward representation and management of mobility.

Recently, the field of data management for moving objects has focused on capturing semantic aspects of moving objects (their locations, trajectories, features, activities, behaviors, etc.). As defined in Antoniou, Corcho, Aberer, Simperl, and Studer (2012), "Semantic data management refers to a range of techniques for

the manipulation and usage of data based on its meaning". According to Lim, Wang, and Wang (2007), thanks to semantic data management, "users can query the data, the domain knowledge, and the knowledge inferred from the data in the same way as querying just relational data". The work presented in Fernández, Arias, Martínez-Prieto, and Gutiérrez (2013) indicates that semantic technologies enable the management of the variety aspect of Big Data. In Bizer, Boncz, Brodie, and Erling (2012) two classes of challenges are identified for Big Data (engineering challenges and semantic challenges), being the semantic challenges related to the problem of "finding and meaningfully combining information that is relevant to your concern". Semantic data management leads to an explicit representation of the meaning of entities and their features, which can later be exploited to provide a more effective and complete query processing and analysis by enabling reasoning, data sharing and interoperability, and flexible solutions that are able to interpret those meanings and act accordingly rather than rely on hard-coded predefined behavior (time-consuming to develop and inflexible in face of unexpected situations that were not anticipated). Several semantic techniques for geo-spatial information and knowledge representation, annotation, querying, and reasoning, have been developed and exploited. However, semantic concepts, methods and techniques should still be better incorporated and applied to the domain of moving object data management and thus to smart location-based services (Sheth & Perry, 2008; Ilarri, Illarramendi, Mena, & Sheth, 2011).

* Corresponding author at: School of Engineering and Architecture, Department of Computer Science and Systems Engineering, Edificio Ada Byron, Maria de Luna, 1, E-50018 Zaragoza, Spain. Tel.: +34 876555262.
   E-mail addresses: silarri@unizar.es (S. Ilarri), dragan.stojanovic@elfak.ni.ac.rs (D. Stojanovic), cyril.ray@ecole-navale.fr (C. Ray).

We argue that, by mixing methods and techniques developed from different fields like *moving object databases* (*MOD*) and the *Semantic Web* –an extension of the Web in which the semantics, or meaning, of the information is explicitly and formally defined– (Berners-Lee, Hendler, & Lassila, 2001; Shadbolt, Berners-Lee, & Hall, 2006), it is possible to enhance the way information about moving objects is managed and the way services are designed (from the modeling, querying, processing, and analysis points of view). As an example, recent proposals claim that linking semantics and Location-Based Services (LBSs) (Schiller & Voisard, 2004), which are services that provide customized information depending on the current location of the user (e.g., vehicle tracking applications, friend-finder applications, location-based emergency services, location-based advertisements, location-based games, etc.), can provide valuable benefits (Ilarri & Illarramendi et al., 2011), such as:

- *Flexible querying*. In a traditional query processing approach there is a certain predefined data **schema** that the user needs to be aware of in order to submit queries asking about the types of entities defined in the schema and their attributes. By using semantic techniques, more flexible and dynamic approaches are possible, based on keyword-based searching (Trillo, Gracia, Espinoza, & Mena, 2007) or intelligent query-answering (Yu & Spaccapietra, 2010), that can exploit information about the context of the user. This can require the use of semantic techniques such as: word sense disambiguation (Navigli, 2009), to interpret correctly the actual meaning of the keywords introduced by the user; automatic classification, as for example if a user asks about "transportation options" this may include different types of entities, such as taxis and buses; semantic spatial search (Park, 2014); and even automatic service discovery and matching.
- *Management of semantic locations and trajectories*. Locations can be expressed at different levels of granularity and abstraction, and users should be able to use the location terminology they require (e.g., cities, provinces, neighborhoods, buildings, rooms, etc.); for this purpose, these symbolic locations could be encoded in ontologies (Ding, Kolari, Ding, & Avancha, 2007; Gruber, 1993; Horrocks, 2008; Uschold & Gruninger, 2004) that explicitly represent the properties of the locations and their relationships. Similarly, applications should consider trajectories of moving objects that characterize their spatiotemporal behavior beyond just a raw sequence of locations, enriching raw trajectory data with higher-level information (semantic annotations) that can later be exploited by applications. Due to its relevance for our proposal, we discuss semantic locations and trajectories in more detail in Section 2.
- *Interoperability among different LBSs and providers*. Mobile users should receive useful information services independently of their current location, which means that there should be a way for different LBSs (potentially developed by different companies and operating in different areas or countries) to cooperate, by sharing and exchanging data. This requires the use of semantic techniques (semantic matching to bridge the existing differences in the data schemas) and/or standards (common data schemas) that guarantee that the data exchanged are interpreted unambiguously.
- *Protection of personal location information*. The location of a mobile user represents very sensitive information that should be kept as private as possible and used only for the requirements of the service that the user would like to access (Krumm, 2009). However, it has been shown that the simple obfuscation of the location data sent to a server (introduction of a deliberate imprecision) is not enough, as an attacker could correlate the coarse location provided with background

knowledge to try to identify the location with higher precision. So, semantic techniques are needed to prevent those potential attacks, as proposed in the semantic-aware PROBE framework (Damiani, Bertino, & Silvestri, 2010). The privacy issue is mentioned several times along this paper, and particularly in Section 4.5.
- *Reasoning in complex and dynamic contexts*. By effectively benefiting from the explicit representation of the semantics of entities, it is possible to provide intelligent LBSs, that reason over the available information (e.g., the location of the user, the user preferences based on a user profile, etc.) and take smart decisions adapted to the current context of the user.

The ideal is to be able to develop and provide smart location-based and mobility services that understand the user requests and interaction, implicitly based on the semantics of mobility and contextual information, and know how to behave and adapt in dynamic and unexpected situations. For example, users may issue the following information requests to service providers: *"What are the interesting places in the city this evening and what is the best schedule for visiting them?", "What is currently the best evacuation path for people with disabilities?", "Will there be a traffic jam on my trip to the meeting (and, if so, how to avoid it)?"*. Other proposals, such as Viktoratos, Tsadiras, and Bassiliades (2014), that focuses on Location Based Social Networking Services (LBSNS), have also emphasized the interest of semantic technologies to represent physical entities and their associations, enable interoperability among heterogeneous systems and knowledge sharing, and provide a common ground for reuse and future extension.

The creation and exploitation of the semantics of the users' mobility would be a key issue to develop systems that dynamically recommend customized and useful information and services to mobile users and that proactively provide notifications regarding events in their surroundings (Ardle, Petit, Ray, & Claramunt, 2012). To support such requests, a mobile user may need to send his/her current location, (partial) trajectory, personal context, profile and preferences, and any other information necessary for the service providers. However, mobile users do not want to reveal lots of personal and trajectory related information to service providers, even if this could enable them to detect the semantics of their locations and trajectories and intelligently satisfy their requests. Moreover, the continuous flow of location/trajectory data originated from moving objects, along with personal and contextual data, user profiles and social interaction, would result in very high demands for the transfer, storage, processing, querying, and analysis of big mobility data.

Pervasive/ubiquitous computing generally considers huge populations of moving objects of different kinds (people, vehicles, assets, animals, etc.) equipped with mobile devices (smartphones, wearable devices, smart sensors) with increasing computing, communication and sensing capabilities, that form a dynamic computing and communication infrastructure connected by different types of both centralized and mobile ad hoc networks (*MANETs*) (Giordano, 2002).

To provide effectiveness and scalability in managing the semantics of moving objects, a distributed infrastructure is needed, in which each moving object collects, stores, processes and analyzes the semantics of its own mobility data/information. We envision a scenario in which moving objects exchange various categories of information and knowledge (e.g., regarding the location, personal and social status, vehicle condition in case the moving object is a vehicle, activity, behavior, environment and traffic conditions, air pollution, etc.) with others, but also share such information/ knowledge through geo-social networks (Shankar, Huang, Castro, Nath, & Iftode, 2012; Vicente, Freni, Bettini, & Jensen, 2011), social media services, and geospatial information services, enabling a

broad range of applications and services, such as urban dynamic mining, public safety, and environment monitoring. The exchange of sensor and user-generated data and the provision of Volunteered Geographic Information (VGI) (Sui, Elwood, & Goodchild, 2012) emphasize a collaborative participatory sensing and crowdsourcing environment (Reddy et al., 2009; Wirz et al., 2013) that is expected to play a key role in this context.

Incorporating semantics in such an *Internet of Moving Things* should provide valuable semantic information and services to mobile users to support smarter mobility in dynamic environments. However, numerous research challenges arise from the application of the semantic paradigm to the management of moving objects, which are summarized in Fig. 1 and described in Section 4 in more detail. Overall, the key contributions of this paper can be summarized as follows:

- We review the related work and the state-of-the-art in the modeling and representation of semantically-enriched locations and trajectories of moving objects. We provide a good coverage of representative efforts, in order to justify the interest of our proposal and identify existing difficulties, but it is not our goal to perform a comprehensive and detailed survey-like study.
- We present the key research challenges of the semantic management of moving objects, that should lead to smarter mobility management in plenty of moving object application domains.
- We propose SemanticMOVE, a novel and scalable reference framework for the semantic management of moving objects, capable of tackling the identified challenges. It incorporates distributed storage, processing, reasoning, analysis, and mining of semantic mobility data.

The rest of this paper is organized as follows. Section 2 presents related work in semantic modeling and representation of moving object locations and trajectories. Section 3 presents two illustrative use case scenarios where the semantic management of moving objects would be beneficial. Section 4 reviews the research challenges to be faced. Section 5 proposes the distributed framework SemanticMOVE for the semantic management of moving objects. Finally, Section 6 provides some conclusions.

2. From raw positioning to semantic locations and trajectories

Various positioning methods and techniques provide the determination of a location represented as a spatio-temporal point (x, y, t) in an appropriate geographic reference system. In open space (outdoors), mobile devices/users are most commonly and accurately positioned by using GNSS (Global Navigation Satellite System) technologies like GPS (Global Positioning System). In indoor environments, accurate positioning is more difficult to obtain: radio technologies such as Bluetooth, RFID (Radio Frequency Identification), NFC (Near Field Communication), and Wi-Fi, are widely used, but infrared and ultrasound technologies or inertial sensors are also of great interest. To provide smooth positioning and navigation across outdoor/indoor spaces, the integration and seamless handover between various positioning systems and technologies is required. Positioning methods provide a stream of location data at different accuracy and sampling rate and determined in a corresponding geometric/geographic space. These locations of moving objects form a sequence of timestamped points delivered to data management systems, for appropriate storage, processing, analysis, and use in a variety of application domains such as animal/people tracking, emergency management, traffic monitoring, location-based social networks, etc.

To provide a semantic interpretation of movement data, recent research has focused on the semantic annotation of locations and trajectories. Related work in the domain of semantics for moving objects mostly considers either the conceptual modeling and analysis of semantic trajectories or the corresponding representation and implementation tied to specific Database Management Systems (DBMSs). However, there is a significant need of methods
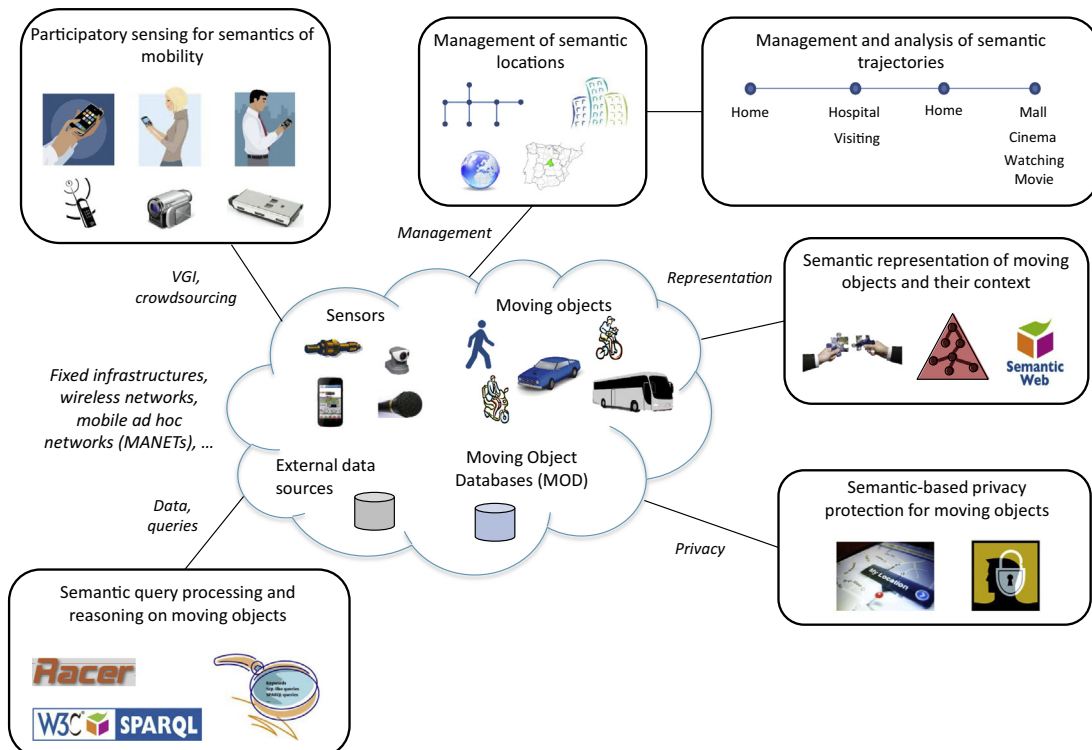

Fig. 1. General overview of semantic mobility data management challenges.

and frameworks for a comprehensive and generic semantic-based representation, management, analysis, and reasoning for moving objects in big mobility data applications.

In this section, we overview some recent work related to the semantic enrichment of locations and trajectories.

2.1. Modeling and management of semantic locations

An essential feature of moving objects is their continuously-changing location. Locations of moving objects can be represented at different levels of granularity, not only at the finest geographic location granularity (Hornsby & Egenhofer, 2002). In other words, one can see the location of a moving object from different semantic perspectives, thus supporting different concepts of location (e.g., GPS coordinates, the city where the moving object is located, the building where it is, the room, etc.), and it should be possible to specify the semantic location terminology that is needed or convenient for each specific query and information need. The concept of *location granule* has been proposed in Ilarri, Bobed, and Mena (2011) to denote these semantic location requirements and describe their impact on the query processing. In other papers, the term *place* is used, for example, to represent human daily activities as a sequence of locations that comply with a certain upper bound distance and a lower bound duration (time-based clustering of location measurements) (Kang, Welbourne, Stewart, & Borriello, 2004) or as a sequence of place visits (Do & Gatica-Perez, 2014); according to Do and Gatica-Perez (2014), the concept of place is "key for studying individual mobility patterns".

Moreover, these semantic locations could also have attached other semantic information that could be exploited automatically during the query processing. For that reason, it is interesting to link them to concepts in an ontology (Ding et al., 2007; Gruber, 1993; Horrocks, 2008; Uschold & Gruninger, 2004). However, how to efficiently manage and effectively exploit these *semantic location granules* (Bernad, Bobed, Mena, & Ilarri, 2013) is still an open problem.

2.2. Modeling and management of semantic trajectories

Trajectories should not be seen exclusively as a sequence of geographic coordinates, but as something that has a meaning, which enables a suitable interpretation of the movement performed. A raw trajectory can be enriched with related contextual information (*semantic annotations*, such as background geographic information, e.g., POIs –Points of Interest–, or other context knowledge). This enrichment can be performed either by the users themselves or automatically by using segmentation techniques and exploiting sensor data (e.g., see Do & Gatica-Perez, 2014). In the following, we discuss the main ideas of semantic trajectories and the proposed modeling approaches.

2.2.1. The basics of semantic trajectories: episodes, stops, and moves

A *semantic trajectory* has extra data that usually describe the trajectory using symbolic entities (locations) at which the moving object stops and moves (along specific paths over a road/street network and possibly using different means of transportation). More generally, a *segmentation* of the trajectory in *episodes* can be performed. According to Mountain and Raper (2001), an episode is a discrete time period for which the spatio-temporal behavior is relatively homogeneous. In general, we could say that an episode is a fragment of the trajectory that satisfies a certain predicate and is significant for a given application (Parent et al., 2013). A traditional segmentation is based on whether the object is moving or not (that leads to distinguishing between *stops* and *moves*), but other criteria could be used instead, such as the transportation mode used (useful for transportation planning in a city, for example) or the type of road traversed (a highway, an urban road, a bike path,

etc.). The proposal presented in Wang, Li, Jiang, and Shi (2013), that focuses on semantic trajectories from videos (detection of video events and video event pattern mining), indicates that a representation based only on *stops* and *moves* is weak (unsuitable to represent, for example, a person that passes a shop without stopping) and so also introduces the semantic tags *enter*, *leave*, *begin*, and *end*, along with an algorithm that generates semantic trajectories from trajectory data using all these tags.

It is interesting to note that episodes are naturally annotated with the criterion that defines the episode (*defining annotation*) but possibly also with other *episode annotations* (Parent et al., 2013), such as the closest POI to the corresponding segment, a semantic region (land use, administrative regions, etc.) or a road network (Yan, Parent, Spaccapietra, & Chakraborty, 2010), or other related features (type of place, activities performed there, etc.). Moreover, the trajectory itself may have other annotations that do not apply to any particular segment but to the trajectory as a whole (Spaccapietra & Parent, 2011). For example, Spinsant, Celli, and Renso (2010) propose a rule-based approach to identify the most probable POI associated to each stop in the trajectory and then globally annotates the trajectory (based on the POIs visited) with information about the overall activity that it represents (e.g., tourism, shopping). Finally, according to Parent et al. (2013), it is also possible to annotate individual trajectory positions, although this is not usual because it could generate a large number of annotation repetitions.

Several proposals have focused on the problem of automatically detecting *significant stops* from the raw trajectories. For example, Ashbrook and Starner (2003) try to detect locations that are semantically meaningful for a moving object by clustering GPS data. Krumm and Horvitz (2006) focus on the prediction of destinations and segment a trajectory into subtrajectories according to certain places identified (gaps of at least five minutes in the GPS samples –indicating that the vehicle was not moving– or at least five minutes of speeds below two miles per hour –interpreted as a parked vehicle with the GPS receiver sending slightly different locations due to noise–). Andrienko, Andrienko, and Wrobel (2007) also consider a temporal threshold to identify significant places within a trajectory (the authors show an example using two hours as the temporal threshold). Zheng, Zhang, Ma, Xie, and Ma (2011) tackle the problem of recommending friends and places to mobile users and define *stay points* as groups of consecutive GPS locations that are within a certain distance threshold and above a certain temporal threshold. Palma, Bogorny, Kuijpers, and Alvares (2008) present a clustering-based algorithm (*CB-SMoT*, *Clustering-Based SMoT*) to identify stops and moves based on the idea that interesting places are determined by the parts of a trajectory where the speed is lower than in other parts of the same trajectory, whereas the method *DB-SMoT* (*Direction-Based Stops and Moves of Trajectories*) (Rocha, Oliveira, Alvares, Bogorny, & Times, 2010) consider the variation of the direction to create the clusters that determine the interesting places. Alvares et al. (2007) present an algorithm (*SMoT*, *Stops and Moves of Trajectories*) that considers as stops those candidate places where the object stays for at least a certain amount of time; the candidate places and the time thresholds are defined by the specific application considered, and each candidate place can have a different time threshold associated. Cao, Cong, and Jensen (2010) have also the goal of extracting *semantic locations* (locations that are semantically meaningful to users) from raw GPS data, by applying clustering techniques, considering the patterns of visits to the clusters and using a reversed geocoder (to obtain street addresses from sample points within the clusters) and yellow pages (to obtain the semantics of those sample points). As a final example, in the computing platform for spatio-temporal trajectories presented in Yan et al. (2010) the authors indicate that they have implemented velocity, density,

orientation and time series based algorithms to identify stops and moves (although Yan et al. (2010) only address the velocity-based approach).

Obviously, which stops are significant and which stops are not is usually application-dependent (e.g., see Alvares et al., 2007), but some techniques try to deduce significant stops without the need to specify the potentially-interesting geographic places in advance (e.g., Palma et al., 2008). The work presented in Andrienko, Andrienko, Hurter, Rinzivillo, and Wrobel (2013) considers application scenarios where there is no predefined set of places from which the places of interest can be selected (e.g., flight or maritime traffic scenarios, animal migrations, sociological studies of human behavior, etc.), and proposes an approach based on event clustering to extract significant places in the trajectories. On the other hand, the proposal in Yan, Macedo, Parent, and Spaccapietra (2008) distinguishes between the *geometric component* of a trajectory (sequence of significant stops and moves, obtained through trajectory segmentation and considering what makes a stop significant from the point of view of the application), the *geographic component* (relevant geographic information data –such as POIs– that characterize the trajectory beyond simple spatio-temporal points, giving an explicit meaning to the stops and moves), and the *application domain component* (application objects that are relevant for the meaning of the trajectory, such as events that may have influenced the movement).

As commented before, it is possible to semantically segment a trajectory to identify episodes different from the traditional stops and moves. For example, the approach in Xie, Deng, and Zhou (2009) tackles the problem of extracting activities (e.g., dining out, working, shopping, etc.) from trajectories by exploiting the concepts of *influence* and *influence duration*. The basic idea is that a certain activity can only occur if there is an appropriate nearby POI during a certain amount of time (e.g., having lunch if there is a restaurant used for at least 30 min). Similarly, Spinsant et al. (2010) also highlight that different activities can be characterized by different durations. As a final example, the proposal in Liao, Fox, and Kautz (2005) presents a location-based activity recognition framework based on Relational Markov Networks.

In a comprehensive survey on semantic trajectories presented in Parent et al. (2013), the concept of *semantic trajectory* highlights the importance of representing the semantics behind a trajectory. That paper reviews existing approaches and techniques for enriching trajectories with semantic information, as well as the use of data mining to analyze semantic trajectories.

2.2.2. Some significant approaches to model semantic trajectories

The work presented in Spaccapietra et al. (2008) focuses on the conceptual modeling of semantic trajectories and proposes two different approaches. One solution implies representing the trajectories and their semantic units (*stops*, *moves*, etc.) in the database schema, whereas the second one implies defining a new *Trajectory-Type* data type to hide and abstract the access to the different components of a trajectory. Semantic trajectories are generated by associating the raw trajectory points and segments with background geographic and application-specific information.

SeMiTri (Yan, Chakraborty, Parent, Spaccapietra, & Aberer, 2011) is a framework proposed for the semantic annotation of trajectories. It is based on a conceptual model called *Semantic Trajectory Model*, which considers a semantic trajectory as a sequence of *episodes* that have a particular significance for a given application. Algorithms are provided to annotate trajectories with *semantic regions* (e.g., representing a building, a university campus, etc.), *semantic lines* (road segments along with the transportation mode used to traverse them), and *semantic points* (meaning of the *stop episodes*; e.g., a bar, a restaurant, a shop, etc.). On the other hand,

SeTraStream (Yan, Giatrakos, Katsikaros, Pelekis, & Theodoridis, 2011) focuses on online trajectory annotation and offers real-time capabilities for cleaning, compressing, and segmenting location data.

CONSTAnT (*CONceptual model of Semantic TrAjecTories*) (Bogorny, Renso, de Aquino, de Lucca Siqueira, & Alvares, 2014) is a semantic trajectory conceptual data model, which accounts for the modeling of semantic subtrajectories, semantic trajectory points, geographic places, events, goals/purposes of each trajectory, environments of the corresponding moving objects, and trajectory behaviors. The proposed general model supports the classical concept of *stops* and *moves* but only as a particular case.

Several proposals exploit ontologies (Ding et al., 2007; Gruber, 1993; Horrocks, 2008; Uschold & Gruninger, 2004) to represent knowledge related to mobility. Motivated by the idea that annotating trajectories using a shared vocabulary would be very beneficial, the proposal in Hu et al. (2013) presents an ontology design pattern for semantic trajectories. The solution proposed abstracts itself from specific application domains (as it looks at classes and properties found commonly in semantic trajectories across different domains). The ontology is encoded in the Web Ontology Language OWL (Bao et al., 2012; Bechhofer et al., 2004), allows the integration of different relevant aspects (related geographic information, domain knowledge, and device data), and can be extended by defining subclasses and subroles as needed for specific applications. Wang et al. (2013) also use an ontology and reasoning to match trajectory patterns. Renso, Baglioni, de Macedo, Trasarti, and MonicaWachowicz (2013) use inferencing and a mobility behavior ontology, that conceptualizes the ground concepts of the domain. Baglioni, Macedo, Renso, and Wachowicz (2008) present an approach for the semantic modeling of trajectories and reasoning on them, proposing to encode domain knowledge in an ontology. An ontological framework for semantic trajectory modeling is also proposed in Yan et al. (2008), which presents ontological modules (sub-ontologies) for the geometric, geographic, and application domain knowledge that define a global *semantic trajectory ontology*. A *Moving Object Ontology* (*MOO*) is also presented in Camossi, Villa, and Mazzola (2013), whose purpose (as opposed to work on trajectory segmentation) is to formalize movement patterns and retrieve trajectories that comply with that behavior (in the context of maritime surveillance). Nevertheless, according to Camossi et al. (2013), the scalability of ontology representations for large datasets of movement data is an open issue.

Recently, the concept of *symbolic trajectory* (sequence of temporally-annotated labels) has been proposed as an alternative to semantic trajectories suitable for implementation in a certain spatial database system (Valdés, Damiani, & Güting, 2013). Specifically, the authors have integrated those symbolic trajectories within the *SECONDO* database system (Güting, Behr, & Düntgen, 2010). The associated labels can then be exploited by a pattern language that can be used to match and rewrite symbolic trajectories. In Dodge, Laube, and Weibel (2012), the authors propose a symbolic representation of trajectories based on the identification of similar movement parameters (e.g., speed, acceleration, direction).

2.2.3. Exploitation of semantic trajectories

Semantic trajectories can be exploited for a variety of purposes. First, as they incorporate annotations that explicitly represent the knowledge behind the trajectories, they enable understandability and sharing. That knowledge can also be used to perform advanced data mining tasks that would be difficult or impossible to perform on the raw trajectories (see Section 4.4 for more details). Overall, the semantic annotations can be exploited for a variety of purposes. For example, Richter, Schmid, and Laube (2012) use

the semantics of a trajectory to compress it without much information loss. As another example, although semantic trajectories can expose more sensitive information, it is possible to consider their encoded knowledge to prevent privacy attacks (see Section 4.5).

Finally, it is interesting to mention the proposal presented in Su, Zheng, Zheng, Huang, and Zhou (2014), which highlights some disadvantages of semantic trajectories (emphasis on moving paths –spatial information– but not on moving behaviors –temporal information–, difficulty to automatically identify the interesting parts in the trajectory, and difficulty to communicate and store them). Motivated by these shortcomings, they present the system STMaker, which obtains a high-level summary of a trajectory in textual format. Although we do not believe that the text summary can replace a full-fledged semantic model of the trajectory, it represents another possible way to exploit the extracted information for specific applications (e.g., the textual summary could be shown to a final user).

2.2.4. Importance of the environment of the trajectories

To conclude this section, it is interesting to mention that, besides the locations and the trajectories, the modeling of the spaces where moving objects evolve is a key semantic feature (Afyouni, Ray, & Claramunt, 2012). For example, the data model proposed in Xu and Güting (2013) supports modeling multiple *infrastructures* (public transportation networks, free space, road networks, and indoor scenarios), each representing a certain environment for the moving objects and defining the possible places where the moving object could be located (e.g., roads and streets for the case of a road network), thus enabling queries considering scenarios with different transportation modes and environments. In Hu et al. (2013), transportation networks are emphasized as an important type of geographic information used to make sense of trajectories. Several proposals related to the field of moving objects have emphasized the importance of considering the underlying transport network for modeling and querying (e.g., Cao & Wolfson, 2005; Ding & Deng, 2011; Güting, de Almeida, & Ding, 2006; Vazirgiannis & Wolfson, 2001). A semantic representation model that identifies significant points within a network-constrained trajectory is presented in Li, Claramunt, Ray, and Lin (2006). Semantic trajectories where the movements of the objects are constrained to a network (e.g., cars, trains) are called *semantic map-matched trajectories* in Parent et al. (2013).

2.3. Beyond traditional semantic-based mobility

Although semantic trajectories have been extensively researched at the conceptual modeling level, and conceptual data models have been developed to represent semantic trajectories, we still lack a generic mobility data management framework.

In the related work, server-based solutions for the semantic enrichment of trajectories are proposed, which follow an off-line and bottom-up approach. They assume that raw spatio-temporal data originated from moving objects are collected, processed and semantically enriched at the central server(s). For example, Pelekis, Theodoridis, and Janssens (2014) envision a framework consisting of three layers: (1) a traditional MOD at the bottom layer, which is in charge of the raw mobility data management; (2) a Semantic Mobility Database (SMD) at the middle layer, that provides novel data types, indexing methods, and operators extending MOD query languages for querying and analyzing mobility data from a semantic perspective; and (3) the application interface at the top layer, providing users with querying and analysis functionality on either the MOD or the SMD that lies below.

Similarly, the *SeMiTri* system architecture (Yan, & Chakraborty, et al., 2011), referenced before in Section 2.2.2, consists of a *stop/*

move computation part that processes the raw GPS records to produce the output trajectory in the form of stops and moves. The Semantic Annotation part of the architecture comprises three annotation layers, for region, line and point annotations, which are combined to produce the annotated trajectory. A Semantic Trajectory Analytics Layer encapsulates methodologies that compute statistics about the trajectories (e.g., trajectory patterns) and stores them as aggregated information in the Semantic Trajectory Store to be accessible by applications.

The server-based approach considered in the existing work means that the moving objects are supposed to continuously send a stream of raw location data to a trajectory server (or several servers, in case a distributed fixed infrastructure of servers in charge of different geographic areas is considered) that constructs and stores the trajectories of the objects in an appropriate MOD, for example SECONDO (Güting et al., 2010) or HERMES (Pelekis, Frentzos, Giatrakos, & Theodoridis, 2008), and performs the semantic enrichment of such trajectories using geospatial information sources (e.g., information about the road network, information about the places visited, geo-sensor networks, etc.). In order to enrich the trajectories with the personal context of the user (e.g., user activities, physiological and emotional status, user profile and preferences), such data need to be continuously sent by the moving objects to the server in its raw form to be integrated and included in the semantic trajectory enrichment process at the server itself. The main shortcomings of this kind of approach are the following:

- Maintaining raw trajectory data of a huge population of moving objects along with rich personal/environmental/social context results in very high demands for network bandwidth, as well as for the storage, processing, querying and analysis of such big mobility data, necessary for the semantic enrichment of trajectories and their exploitation.
- The mobile users do not have control over the semantic enrichment process, i.e., which data are included and how the semantic enrichment is performed.
- A pure server-based approach cannot benefit from the processing and storage capabilities of the mobile devices.
- The privacy of the mobile users could be compromised at the server, where the complete semantics of the users' mobility are collected, analyzed and stored, so the privacy needs to be preserved according to the preferences supplied by the users.

Beyond the current state of the art, a semantic mobility data management framework should provide an efficient and comprehensive storage, querying, processing, analysis and mining of semantic mobility data. It should also fully set up the concept of ubiquitous systems (Weiser, 1991) in a context of mobility that makes it usable in real moving object applications aiming at a better understanding and exploitation of mobility. Moreover, more research is needed to determine how mobile applications can benefit from semantic trajectories (through querying, reasoning and analysis), according to the context of the user. Interoperability situations where there is a need to exchange and understand semantic mobility information collected and defined by others is also an issue to address.

In the next section, we present use case scenarios to show the necessity and the potential of the semantic management of moving objects and mobility data. Several research challenges must be addressed and integrated to provide an efficient and effective semantic management of mobility beyond the state-of-the-art semantic trajectories to make these scenarios a reality. We discuss those challenges in Section 4.

3. Use case scenarios

We first present use case scenarios that illustrate the need for semantic location-based services based on a comprehensive management of the semantics of moving objects.

3.1. Semantics in urban mobility

Alice travels to another city with the assistance of the SemanticMOVE mobile "travel" application. The application analyzes her trajectory, detects a mode of transport (e.g., see Biljecki, Ledoux, & van Oosterom, 2013; Yu, Yu, Wang, Lin, & Chang, 2014), which in this case is a train, and, by exploring her calendar, recent communications and social network activities, recognizes that she intends to go to a meeting. Considering her privacy profile, the application obtains the best possible means of public transportation from the train station to the meeting location. The mobile "travel" application does not need to continuously send her raw locations/trajectory data, nor additional information regarding her current context, social activities and calendar to a service provider, as this information is analyzed locally. So, the application detects the semantics of her mobility from data obtained from the mobile device itself, eventually accessing external web services (e.g., to correlate the geographic location of the user with potential places of interest), and sends only high-level semantic requests to a service provider when needed.

At the same time, Lucia drives a car to the meeting and her mobile "navigation" application receives navigation instructions, notifications, and recommendations, based on the semantically-enriched mobility data stored, processed and analyzed on her device, privacy preferences, and up-to-date traffic information obtained from a third-party service. By exploring the semantic trajectories of other mobile users (friends), different software components identify that the current trajectory of Lucia allows her to pick up Alice at the train station. The navigation application proposes to share its trajectory with Alice and send a "meet-together" request.

After the meeting, Alice and Lucia meet Bob and Miguel, and they decide to go out all together. By collecting and analyzing semantic trajectories and mobility information of other moving objects in their surroundings performing "going out" activities (having dinner, going to the cinema, having a drink, etc.), their mobile "fun and culture" applications exchange information with others and learn about a popular music festival being held in the city. Their mobile applications do not need to send their locations and trajectories to some service provider, along with all personal/environmental/social context information necessary to semantically enrich their trajectories. Instead, they generate semantics of their mobility data at their own devices and exchange semantic information and knowledge according to their privacy preferences with mobile applications/devices in the vicinity, in order to exploit semantic collaboration and crowdsourcing.

3.2. Semantics of mobility in emergency management

When Alice, Lucia, Bob, and Miguel arrive at the festival, they receive the recommendation to use a mobile "festival" application, that helps them navigate through the festival facilities. Their "festival" applications obtain information from the festival service about interesting events and schedules, receive appropriate notifications and recommendations, and exchange such information with other visitors' mobile applications.

In case of emergency (e.g., a fire), their mobile devices act as a distributed semantic sensor platform providing the detection of the semantic movement of the visitors: their activities and behavior, user-generated content (messages, tags, comments, photos),

the parameters of their environment (e.g., high level of CO_2), and their physiological (e.g., injured, tired, dizzy) and emotional status (e.g., frightened, in panic, calm). Being chronically ill, Bob carries a wearable wireless sensor device for heart rate monitoring connected to his smartphone and managed by a mobile "be healthy" application. His device does not have air pollution sensors but his mobile application can obtain such information from applications/devices in his surroundings. During this emergency situation, Bob becomes very nervous, so the application detects a suspicious heart rate and a worrying overall physiological status, as well as high levels of CO_2 and NO_x obtained from the mobile applications/devices in the vicinity, and alerts him. The application also sends a request to the festival emergency service for the best path to the closest ambulance. His mobile application dynamically helps Bob navigate through the festival site, while continuously tracking his physiological status and the environment conditions. Bob receives appropriate notifications and recommendations from the service reflecting the current crowd behavior and the rescue process.

The visitors' devices exchange information, semantic concepts, and knowledge, regarding navigation and evacuation assistance, with others moving in their surroundings. Besides, they can send processed and semantically-enriched, but privacy-preserved, mobility data to a "festival service" (this particular service might be centralized) for aggregated/integrated processing, querying, reasoning, and analysis. This enables the real-time detection of group mobility behavior (flocks, crowds, etc.), a better understanding of the emergency situation, and efficient evacuation decisions. Moreover, other available data sources, such as geospatial information services, social networks, semantic sensor networks, and web services (Alonso, Casati, Kuno, & Machiraju, 2004; Alonso, 2002; McIlraith, Son, & Zeng, 2001; Payne & Lassila, 2004), can be exploited.

Thanks to this collaborative approach, Alice, Lucia, Bob, Miguel, and others, receive timely notifications with optimal directions according to their own mobility, situation, and behavior, as well as the overall group behavior and movements. Besides, pictures and videos recorded by the participating mobile devices could be analyzed and correlated with their trajectories to try to identify the cause of the fire.

4. Research challenges towards better semantics of mobility

The study of the semantics of moving objects, their trajectories, and mobility data, is still at its early stage and there are numerous research challenges that need to be tackled in the coming years. In this section, we review those challenges.

4.1. Participatory sensing for semantics of mobility

With the proliferation of mobile devices embedding sensing capabilities, users with their mobile devices have become an important source of sensor data (Lane et al., 2010; Lane, 2012). Different from the traditional sensor motes in wireless sensor networks, these mobile devices are equipped with an increasing number of built-in sensors: GPS, microphone, camera, ambient light sensor, accelerometer, gyroscope, compass, proximity sensor, and even temperature and humidity sensors, possessing also increasing computing and communication capabilities. Moreover, through the user interaction with a mobile device, so-called virtual and social sensors can be defined to detect active applications, user activities, social network connections, privacy preferences, etc., but also to provide user-generated content (videos, photos, sounds, texts, speech messages) referenced in space and time. All these sensors enable crowdsourcing approaches for sharing data about

the environment among users (Dobre & Xhafa, 2014; Reddy et al., 2009; Wirz et al., 2013; Zaslavsky, Perera, & Georgakopoulos, 2013).

Advances in wearable technologies, such as the popular *Google Glasses*[1] or *Samsung Galaxy Gear*,[2] will boost in the near future data capture and sharing. The development of projects such as *Sensordrone*[3] and *uRing*[4] also emphasizes the interest in adding even more sensors to measure air pollution, personal health parameters, and the emotional and physiological status of users.[5] As another example, according to Fleming (2013), "Today's luxury cars have more than 100 sensors per vehicle". These sensor data could be shared and exploited in a Vehicular Ad Hoc Network (VANET) (Hartenstein & Laberteaux, 2008; Olariu & Weigle, 2009) or in a hybrid network with both vehicles and mobile users (Liu, Liu, Cao, Chen, & Lou, 2010).

Thanks to these capabilities, each mobile device can continuously capture and process spatio-temporal sensor data that represent low-level context information and that are semantically annotated (Calbimonte, Yan, Jeung, Corcho, & Aberer, 2012). The application of analysis, reasoning and interpretation techniques on the low-level context augments and annotates a moving object trajectory with high-level semantics that describe its mobility through a certain environment and represent the moving context/situation of such an object.

The research in this area focuses on the representation, filtering, processing and analysis of large volumes of mobile sensor data that describe the user movement and activities (walking, running, moving in a wheelchair, driving, etc.), interaction (talking, searching on the Web, tagging a photo, connecting with friends, etc.), environment conditions (e.g., bad weather conditions, traffic congestion, crowded area), and preferences (Predic, Yan, Eberle, Stojanovic, & Aberer, 2013). Semantic techniques are being proposed for the annotation of sensor measurements, in order to enhance the interoperability and provide useful contextual information (Sheth, Henson, & Sahoo, 2008; Calbimonte, Jeung, Corcho, & Aberer, 2012). The analysis and mining of semantic mobility data and their context should provide insights into important features of the movement, the situation, and behavior of the moving object, as well as the prediction of its future movements.

4.2. Semantic representation of moving objects

Another important challenge is how to represent the semantic information associated to a moving object in a way that enables interoperability, querying, analysis, and reasoning. The idea is to manage knowledge about the types of moving objects, the different profiles that they may show to others, and their mobility features, in such a way that this knowledge can be exploited later (e.g., for analysis or for query processing). From a representational point of view, there should be no distinction between mobile users and other moving objects.

We argue that the semantic information attached to a trajectory should be represented at different levels of detail and it should be possible to easily commute from one to another based on the query requirements and existing privacy constraints. Thus, subsequences of raw trajectories can be semantically annotated at different levels of granularity. For example, at the higher level of granularity a user moves from one city to another, where cities represent *stops* along highways, which represent *moves*. At an intermediate level, *stops*

are represented as POIs within the city, and *moves* as street segments. Finally, at a smaller granularity level the user moves between buildings, floors, offices, halls, etc., along the corridors, stairs, elevators, etc.

The concept of semantic granularity can be applied to other semantic attributes of a moving object. For example, at the coarsest level user physical activities could be described as *steady* or *active*; at the finer level the *steady* state can be *sitting*, *standing*, or *lying down*, while *active* can be *walking*, *running*, *driving*, etc. An analogous granularity hierarchy and decomposition can be defined for other semantic attributes that represent a higher-level context of a moving object and that are generated through a semantic annotation of raw contextual data.

Thus, the semantics of the mobility of a moving object can be represented along various mobility dimensions and at different levels of granularity, as illustrated in Fig. 2. In the figure, we can see that the same geographic trajectory could be represented as a semantic trajectory at different abstraction levels, depending on whether the focus is on the different means of transport used, the activity performed, the exposition to pollution, etc. For example, during the trajectory shown in the figure, the spatial geographic area of the park corresponds to different semantic trajectory segments: "Park" (according to the stop/move trajectory), "By foot" (according to the means of transport), "low" (regarding the CO_2 level), "Jogging" (regarding the general activity performed by the user), and "Heart rate alert" (according to the information provided by the physiological sensors worn by the user). The figure also shows that some representation levels may be applicable only in some cases (e.g., a specific symbolic trajectory while the user is indoors –fourth row of trajectories in the figure–) and that some segments of a representation level may be unknown for some time intervals (e.g., as it is the case of the user's mood in the example).

The context, that semantically enriches the locations and trajectories of moving objects, can be modeled and represented in different ways, such as using key-value or markup schemes, graphic and object-based techniques, as well as using ontologies. According to several surveys (Baldauf, Dustdar, & Rosenberg, 2007; Bettini et al., 2010; Perera, Zaslavsky, Christen, & Georgakopoulos, 2014), ontologies are the preferred technique for managing context and reasoning over it in context-aware systems, having in mind also their shortcomings regarding high storage and performance requirements when the amount of semantic data increases.

As the semantic annotation and enrichment of moving objects and their trajectories is performed in a semantic way, the use of ontologies definitely can play a key role, as well as semantic web standards (e.g., RDFS, OWL) and related technologies, such as reasoners (Dentler, Cornet, ten Teije, & de Keizer, 2011; Mishra & Kumar, 2011). Services that a moving object can access, or can find interesting at a certain moment, depend on the continuously changing context, i.e., a *contextual trajectory*; the relevant services may change or behave differently according to the changes in the context. Context-aware services (Baldauf et al., 2007; yi Hong, ho Suh, & Kim, 2009) and semantic location-based services (Ilarri & Illarramendi et al., 2011) should also fully exploit contextual dimensions such as the user-centered context (e.g., user profile, preferences, user's physical/cognitive capabilities), the environmental context (e.g., location, light), the temporal context (e.g., day of the week, season, time period), and the context of execution (e.g., network connectivity, nearby resources). The use of semantic trajectories can help in the representation of the moving context, or *contextual trajectory*, as locations and times in a semantic trajectory usually define the core of a context model. Moreover, knowledge about the environment itself may also be relevant in certain situations. For example, for objects moving within a building the layout of the building can be relevant for indoors

---

[1] http://www.google.com/glass/start/.
[2] http://www.samsung.com/us/mobile/wearable-tech/SM-V7000ZKAXAR.
[3] http://sensorcon.com/products/sensordrone-multisensor-tool/sensordrone.
[4] http://usenns.com/en/.
[5] http://mobihealthnews.com/11224/by-2016-80m-wearable-wireless-fitness-sensors/.
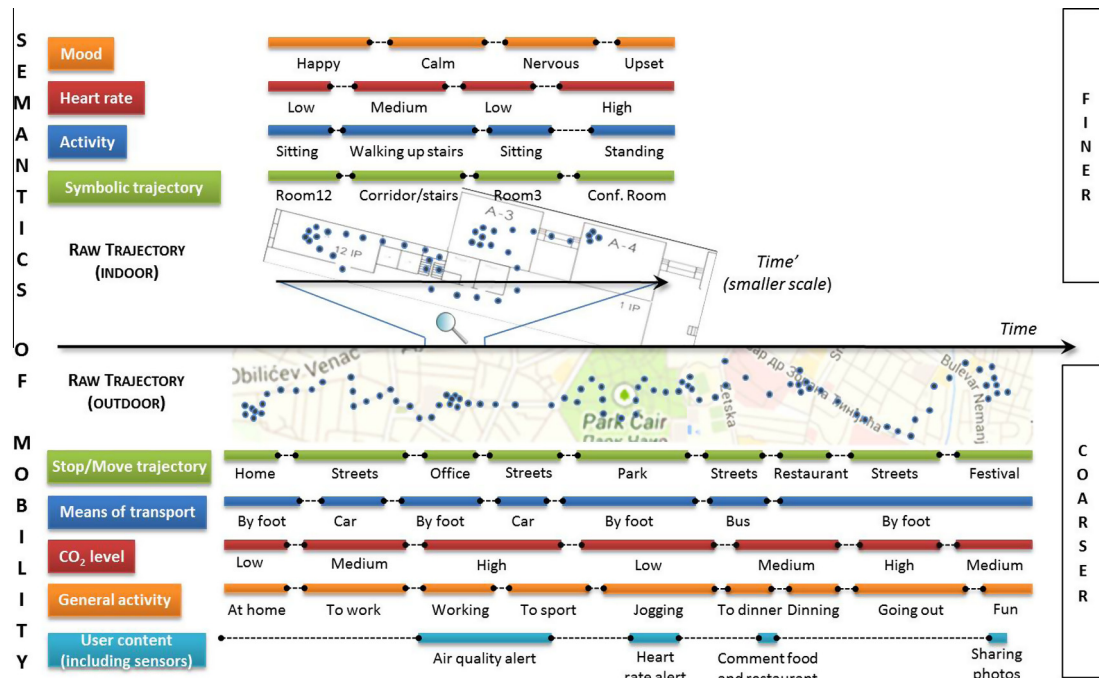
Fig. 2. Semantics of a moving object represented at different levels of granularity.

query processing, whereas in the case of objects moving freely in the geographic space considering a geographic map as part of the context could be useful; it is possible that even a different notion of distance should be applied in each case (network-based distance vs. Euclidean distance). As an example, the work presented in Ben, Qin, and Wang (2014) focuses on indoor moving objects and emphasizes that existing proposals consider the locations of the moving objects but ignore the semantics of the objects themselves, which motivates their proposal of a new *indoor semantic-based index*.

Proposed solutions should not rely on centralized knowledge repositories and rather favor distributed and fragmented knowledge that can be integrated dynamically (in an ad hoc manner). The challenge is how to integrate and combine all the semantic aspects and manage them as required in a uniform and efficient way. Techniques related to the fragmentation and modularization of ontologies (d'Aquin, Sabou, & Motta, 2006; Jiménez-Ruiz, Grau, Sattler, Schneider, & Berlanga, 2008) and techniques for the exchange and integration of pieces of knowledge (Calvanese, Giacomo, & Lenzerini, 2002; Choi, Song, & Han, 2006) are relevant here.

4.3. Semantic query processing and reasoning

Even with all the semantic elements in place, a key issue is how to exploit the semantic information in the query processing in a way that is both efficient and effective. This involves using reasoning in a truly useful way that allows inferring useful facts that have not been explicitly asserted and exploiting them during the query processing, leading to a special type of location-dependent queries (Ilarri, Mena, & Illarramendi, 2010) with semantic awareness. Besides, users should be able to submit queries in a flexible way, for example by expressing keyword-based queries (Trillo et al., 2007), without requiring them to know any kind of schema information or single knowledge base.

To accomplish this challenge, a variety of semantic data may need to be managed, regarding locations, trajectories, moving object representation, and sensor data. For example, reasoning

about the context (Bettini et al., 2010) could be a basic building block to enable semantic-based data management on moving objects. Reasoning about mobility data (locations, trajectories, sensor data, etc.) provides the generation of high-level motion context information based on raw sensor data and lower-level context through inference methods and sensor fusion (Bikakis, Patkos, Antoniou, & Plexousakis, 2008; Castelli, Mamei, Rosi, & Zambonelli, 2009). Ontology reasoning over the ontological representation of mobility data brings new and higher-level semantics to user motion, based on semantic reasoning engines, such as Pellet[6] and HermiT.[7]

As mentioned in Section 4.2, a reasoning process will also need to take place for ontology alignment (Calvanese et al., 2002; Choi et al., 2006) when objects exchange pieces of knowledge (i.e., fragments of ontologies). Moreover, in a distributed framework there should be support to perform some reasoning on mobile devices with limited storage, processing and communication resources as well as limited battery life (Bobed, Bobillo, Yus, Esteban, & Mena, 2014; Motik, Horrocks, & Kim, 2012; Ruta, Scioscia, Sciascio, Gramegna, & Loseto, 2012; Ruta, Scioscia, Loseto, Gramegna, & Sciascio, 2012; Patton & McGuinness, 2014, in press; Yus, Bobed, Esteban, Bobillo, & Mena, 2013). Finally, reasoning over sensor data in the *Semantic Sensor Web* is another hot topic of research (Barbieri, Braga, Ceri, Valle, & Grossniklaus, 2010; Valle, Ceri, van Harmelen, & Fensel, 2009).

4.4. Analysis and mining of semantic mobility data

The research in this area aims at providing an appropriate conceptual modeling and representation of mobility, but more importantly the analysis of massive amounts of mobility data at different levels of granularity. Traditional tasks for knowledge extraction from mobility data or spatio-temporal data mining (Nanni, Kuijpers, Koerner, May, & Pedreschi, 2008) include: searching similar trajectories or *clustering* trajectories (e.g., see Pelekis et al.

---

[6] http://clarkparsia.com/pellet.
[7] http://hermit-reasoner.com.

(2012) or Lee, Han, & Whang (2007) –this last work supports the discovery of common subtrajectories–), based on some defined similarity function that groups trajectories based on certain properties (e.g., space and time, speeds, accelerations, directions); *classifying* trajectories into predefined classes (e.g., see Lee, Han, Li, & Gonzalez, 2008); *discovering patterns* (e.g., see Giannotti, Nanni, Pinelli, & Pedreschi, 2007; Laube, 2009; Wood & Galton, 2009), such as sequences of locations or group behavior (e.g., objects moving together); and identifying *regions of interest (ROIs)* from a set of trajectories, based on the idea that a region is interesting if it contains a large number of moving objects that remain inside for at least a certain amount of time (e.g., see Giannotti et al., 2007; Uddin, Ravishankar, & Tsotras, 2011).

Mining raw data about moving objects (e.g., see Andrienko, Andrienko, & Heurich, 2011; Bogorny, Heuser, & Alvares, 2010) can be useful to extract information that can later be used to build semantic trajectories. For example, in some situations the task of automatic place labeling can be considered as a multi-class classification task (Do & Gatica-Perez, 2014). Moreover, the aggregation and clustering of trajectories (e.g., see Dodge et al., 2012) are of importance to understand the meaning behind the trajectories and try to correlate them with people's trends, habits and behavior patterns, or to understand the behavior of crowds (Wirz et al., 2013). According to Pejovic and Musolesi (2013), "As devices become increasingly intelligent, their capabilities evolve beyond inferring context to anticipating it." Machine learning methods and techniques play a key role here.

On the other hand, it is also possible to analyze semantic trajectories to extract knowledge about their features, such as the behavioral patterns of the moving objects (e.g., see Laube, 2009; Parent et al., 2013; Wood & Galton, 2009). Examples of behaviors that can be considered for individual trajectories are the *tourist behavior* (trajectories involving accommodation, tourist attractions, and eating places), the *speeding behavior* (vehicles traveling above the speed limit), a *sequence behavior* (characterized by a sequence of component predicates in a certain order), etc. There has also been considerable interest in behaviors that apply to groups of objects (e.g., see Laube, Imfeld, & Weibel, 2005; Nanni et al., 2008; Wood & Galton, 2009). Some of these behaviors are collective, as they characterize groups of trajectories, such as the *meet behavior* (trajectories that end at the same point at the same time instant), the *flock behavior* (trajectories that stay close to each other during some time), etc. Finally, some behaviors apply to individual trajectories but in relation to others, such as the *leadership behavior*.

Mining semantic mobility data that enable the analysis and discovery of semantic trajectory patterns is a topic that has not been extensively considered so far in the literature. Nevertheless, for example, the work presented in Bogorny et al. (2014) has emphasized its importance (e.g., to discover the reasons for a traffic jam or for the detection of frequent semantic patterns). The work presented in Ying et al. (2014) exploits the similarity of semantic trajectories to cluster users for the purpose of recommending appropriate items to them. As indicated in Parent et al. (2013), "Considering the semantics of space [...] and the semantics of time [...] gives more meaning to a behavior" and "the integration of trajectories with contextual and semantic, spatial and temporal, information is vital for the discovery of meaningful behaviors" (e.g., "Going from school to cinema on Wednesday afternoon"). In other words, semantic trajectories enable the extraction of behaviors that would not emerge from raw trajectories. Other proposals, such as that of Alvares et al. (2007), also argue that adding semantic information to trajectories facilitates trajectory data analysis in different application domains. Indeed, as shown in Alvares et al. (2007), certain queries about the moving behavior of objects can only be answered by considering trajectories along with their associated semantic information. Moreover, Hu et al. (2013) emphasize

that semantic annotations are necessary to improve the discovery, reuse and integration of trajectory data from different sources, and that "semantically-enriched trajectories facilitate the discovery of new knowledge, which otherwise may not be easily found". The concept of *semantic behavior* (Parent et al., 2013), which is trajectory behavior whose defining predicate is based not only on pure spatio-temporal criteria but also on other context data, is important and its identification and exploitation is a significant research challenge.

Moreover, the work presented in Parent et al. (2013) indicates the need of new computational models for semantic trajectories when very large trajectory datasets are available. Several approaches could be considered for mining massive mobility data, such as on-line processing using main memory and possible approximations, historical processing that manages and analyzes semantic trajectories already stored, and hybrid approaches that combine both. The amount of data involved in the process may be very high. For example, a high number of trajectories may need to be correlated to identify group behaviors or interactions among a large number of moving objects.

As an example of the interest of data analysis in this context, *WhereNext* (Monreale, Pinelli, Trasarti, & Giannotti, 2009) exploits trajectory pattern mining in order to predict the next location of a moving object. The work developed in Dodge, Weibel, and Lautenschütz (2008) attempts to contribute to the development of a toolbox of data mining algorithms and visual analytic techniques for movement analysis by proposing a taxonomy of movement patterns. The work presented in Bogorny, Kuijpers, and Alvares (2009) proposes a specific query language (the *Semantic Trajectory Data Mining Query Language –ST-DMQL–*) for knowledge discovery from semantic trajectories and for the semantic enrichment of raw trajectories (implemented in *Weka-STPM* (Bogorny, Avancini, de Paula, Kuplich, & Alvares, 2011)). Another similar tool is the *M-Atlas* (*mobility atlas*) system and its associated data mining query language, presented in Giannotti et al. (2011). Finally, it is interesting to indicate that a recent survey on semantic trajectories (Parent et al., 2013), mentioned in several places along this paper, covers data mining techniques to extract knowledge from trajectories: clustering, classification, discovery of common sequences of movements, and identification of objects with a related movement pattern.

The main research issues in this area are related to the difficulty to analyze large volumes of data in real-time, as well as performing such a task distributively on mobile devices. Although there is some work on distributed data mining (e.g., Datta, Bhaduri, Giannella, Kargupta, & Wolff, 2006; Kumar, Kantardzic, & Madden, 2006; Tsoumakas & Vlahavas, 2009), it has not tackled mobility data or mobile peer-to-peer (P2P) scenarios. As indicated in Yu et al. (2014), limiting the power consumption of mobile devices and increasing their performance for data analysis are two critical issues to consider.

4.5. Semantic-based privacy protection

Location privacy is an important concern when dealing with information about mobile users. Significant research efforts, materialized in publications (e.g., see Beresford & Stajano, 2003; Görlach, Heinemann, & Terpstra, 2005; Krumm, 2009) and projects (e.g., the *GeoPKDD –Geographic Privacy-aware Knowledge Discovery and Delivery–* project, http://www.geopkdd.eu/, and the *MODAP – Mobility, Data Mining, and Privacy–* project, http://www.modap. org/), have considered the problem of protecting sensitive location data.

The increasing availability of sensors and other data sources may facilitate privacy attacks, due to the risk of combining data from different sources with background knowledge. In addition,

the semantic management of locations and trajectories may lead to an increase of threats to privacy, since the basic movement data can be augmented with rich information about the user, such as points of interest, preferable routes, his/her context while moving, and also his/her activities, behavior, and future movements, that can be processed and analyzed by a third party to reveal very sensitive user information. As indicated in Parent et al. (2013), "semantic trajectories and privacy clash"; however, this conflict can be mitigated by protecting *sensitive semantic locations* (e.g., hospitals, religious places, etc.).

The basic idea is that introducing noise in the geographic coordinates of the objects whose locations must be protected is not enough. Instead, semantic-based approaches are needed, as mentioned in Damiani, Silvestri, and Bertino (2011) and Parent et al. (2013). *Semantic location cloaking* (Damiani et al., 2010, 2011) is based on the idea that locations can be more or less sensitive depending on the context, mainly the place where the person is located, and therefore the location granularity disclosed can be adjusted accordingly. The semantic location cloaking model was proposed for locations in an unconstrained geographic space, but Yigitoglu, Damiani, Abul, and Silvestri (2012) extended the model for the case of objects moving through road networks. There are also some approaches that focus on the anonymization of published trajectories. For example, Monreale, Trasarti, Pedreschi, Renso, and Bogorny (2011) present a framework for the anonymization of semantic trajectories that replaces the visited places by more general concepts when necessary to protect privacy (e.g., "the Louvre museum" could be replaced by "museum" or simply by "touristic place"); the concept of *c-safety* is used to denote an upper bound for the probability to infer that a certain person has visited a sensitive place.

Overall, the privacy problem is a difficult one. The survey presented in Parent et al. (2013) indicates the need of extending privacy-preserving solutions to behaviors. In Musolesi (2014) the question of whether it is possible to combine large-scale mobile data mining technologies with privacy protection is posed.

Research in this area includes the development of effective methods that exploit background information to detect potential privacy attacks and prevent them.

5. The SemanticMOVE framework

The management of mobility data and its exploitation by semantic location-based services requires adapted computing frameworks, which is also a challenge towards the semantics of mobility. As mentioned in Section 2.3, server-based, off-line and bottom-up approaches for the semantic enrichment of trajectories do not fulfill all the requirements. We postulate that a generic and scalable distributed framework comprising the key components and functional blocks dedicated to a comprehensive semantic management of moving objects is necessary to address the challenges described in the previous section. We propose such a framework, called *SemanticMOVE* (http://webdiis.unizar.es/silarri/SemanticMOVE/), that represents our vision for the semantic management of moving objects. It leverages the increasing sensing, processing, interaction, communication and energetic capabilities of mobile devices in a scalable way.

5.1. Framework design

The proposed framework is based on a distributed architecture comprising software and data storage components running on mobile devices (e.g., smartphones) and servers providing a semantic LBS infrastructure (see Fig. 3). The SemanticMOVE mobile components are organized in three layers:

- *Mobility Data Collection Layer*. The first layer performs the collection and aggregation of mobility data and personal context from an increasing number of sensors integrated in a mobile device or attached to the human body, as well as from virtual and social sensors that reflect user interactions and user-generated multimedia content.
  The collaborative collection of data and its pre-processing to extract semantic information is a challenge that was described in Section 4.1.
- *Semantic (Personal) Mobility Management Layer*. This layer gathers data from the Mobility Data Collection layer and provides a semantic representation of personal mobility data and trajectories, based also on information integrated from available (semantic) geospatial data, the Web, and social media services. It performs semantics injection to mobility data managed and stored at the mobile user device according to semantic concepts and ontologies defined for specific domains. Based on the semantic representation of mobility data, this layer provides storage, query processing, reasoning, and analysis functionalities over the semantic representation of moving objects based on semantic techniques such as ontologies (Ding et al., 2007; Gruber, 1993; Horrocks, 2008; Uschold & Gruninger, 2004) and reasoners (Dentler et al., 2011; Mishra & Kumar, 2011) while employing privacy-preserving mechanisms. This layer may benefit from access to a *Semantic Management of Big Mobility Data and Trajectories* layer provided by existing servers. The development of this layer faces several challenges related to the need to support a semantic representation of moving objects (see Section 4.2), a semantic-based query processing (see Section 4.3), and an analysis of trajectories (see Section 4.4), and at the same time preserving the privacy of the moving objects involved (see Section 4.5).
- *Semantic Mobility Application Layer*. The last layer includes a variety of mobile semantic location-based applications (travel, health, going out, culture, festivals, etc.) that use the Semantic-MOVE API (Application Programming Interface) for accessing semantic management functionalities provided by the Semantic (Personal) Mobility Management Layer. This layer may benefit from access to a *Semantic Location-Based and Mobility Services* layer provided by existing servers.
  Once all the core components of the framework are in place, the challenge is to develop interesting applications that can maximize the benefit of SemanticMOVE. Whereas the existence of a specific "killer application" is unlikely, we believe that the flexibility of a framework like SemanticMOVE would enable a variety of interesting applications and services. A key to achieve this goal relies on an appropriate API design and interoperability mechanisms. In Section 3, we presented two illustrative use case scenarios.

If a user requires a smart mobility service, his/her mobile *semantic mobility application* sends a part of the semantic mobility data stored on the user device to an appropriate *Semantic Location-based and Mobility Service*, which returns explicit responses (reactive behavior) or sends notifications and recommendations of both relevant information and useful services (proactive behavior). As mentioned before, to provide smarter mobility to mobile clients, these semantic mobility services can rely on the SemanticMOVE Server components and an API to access their functionalities for the *Semantic Management of Big Mobility Data and Trajectories*. The SemanticMOVE Server components perform the aggregation of semantic mobility data collected from a large number of mobile users/moving objects, and provide query processing, reasoning, analysis, and mining, over massive mobility data sets. They enable the detection of aggregated mobility patterns and trajectories, collective activities and behavior, as well as complex situations
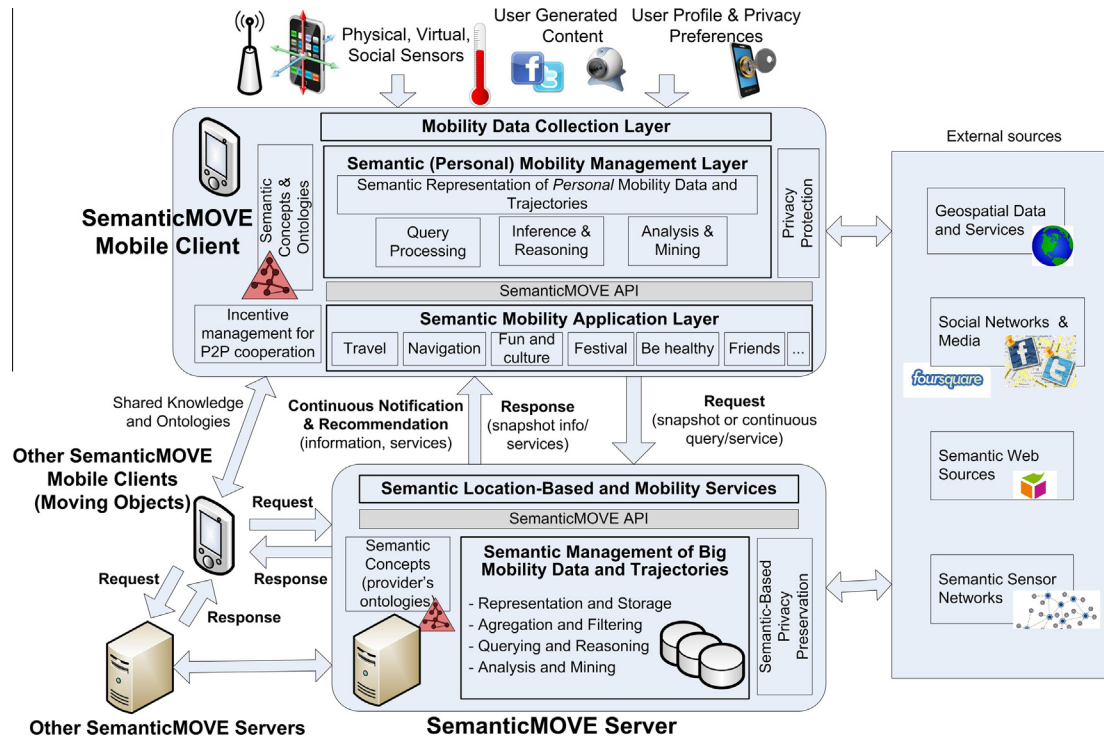
Fig. 3. Architecture of the SemanticMOVE framework.

(e.g., interesting places, traffic congestion, popular city routes, or crowded evacuation paths in an emergency situation). A Semantic-MOVE Server maintains and manages semantic concepts and knowledge related to specific application domains and employs semantic-based privacy preservation mechanisms over a potentially-large number of mobile users. The use of ontologies facilitates the interoperability among components in this heterogeneous environment.

In the rest of this section, we first analyze the potential deployment options of SemanticMOVE. Then, we summarize an analysis of its feasibility.

5.2. Distributed deployment

As opposed to other related work, we envision a quite generic architecture, supporting *a fully distributed scenario for the management of semantics of moving objects*. Thus, each moving object can collect, represent, and analyze its own semantic mobility data and trajectories, and reason over them locally. For this purpose, it can share and exchange semantic mobility data, trajectories, and semantic concepts/knowledge, with other moving objects/ users in the vicinity over ad hoc wireless networks. Our framework includes fixed servers, but not as a requirement but rather as an additional element of an ecosystem where the distributed and ad hoc cooperation among moving objects plays a key role; this cooperation could be encouraged by exploiting some incentive mechanism (Dias, Rodrigues, & Zhou, 2014; Wolfson, Xu, & Sistla, 2004; Xu, Wolfson, & Rishe, 2006), and semantic techniques can be used to enable the interoperability and data sharing between the devices (e.g., see Mandreoli, Martoglia, Penzo, & Sassatelli, 2009). So, we could envision several possible deployments such as those shown in Fig. 4:

- *Purely Centralized* (single node), where a single fixed computer (infrastructure node) plays the role of SemanticMOVE Server, other mobile devices act as SemanticMOVE Clients by communicating information and requesting services from the centralized service provider, and there may also be other mobile devices representing moving objects that provide information but do not request any service (e.g., urban buses that make their location publicly available for the citizens).
- *Distributed with light clients*, where several SemanticMOVE Servers exist. They can host similar or complementary functions, communicate with mobile appliances, and possibly between them. The distribution of responsibilities to the different service providers could be determined, for example, based on geographic criteria: each provider could be in charge of information and requests concerning a certain spatial region (as in approaches like Ilarri, Mena, & Illarramendi, 2006).
- *Distributed hybrid*, where we also introduce more powerful mobile devices that encapsulate both SemanticMOVE Client components and SemanticMOVE Server components. So, not only fixed servers but also those devices can provide advance mobility services to other devices.
- *Purely Distributed*, where only mobile nodes participate, without the assistance of fixed servers. Mobile devices encapsulate a SemanticMOVE Server that others can access, possibly using short-range ad hoc wireless communications.

Distributed architectures provide key benefits for the management and analysis of large amount of data, such as (Tsoumakas & Vlahavas, 2009): reduced storage cost (in comparison with a centralized approach that would need a very large storage space to keep all the data in a single place), reduced communication cost (as large volumes of data do not need to be communicated to a centralized repository), reduced computational cost (the data distribution can be exploited to perform simpler tasks in parallel), and better protection for private and sensitive data (as opposed to a centralized data collection approach, each participant in the distributed environment can keep control over its own data).

The use case scenarios presented in Section 3 illustrate how semantic mobility data can be exploited to provide dynamic
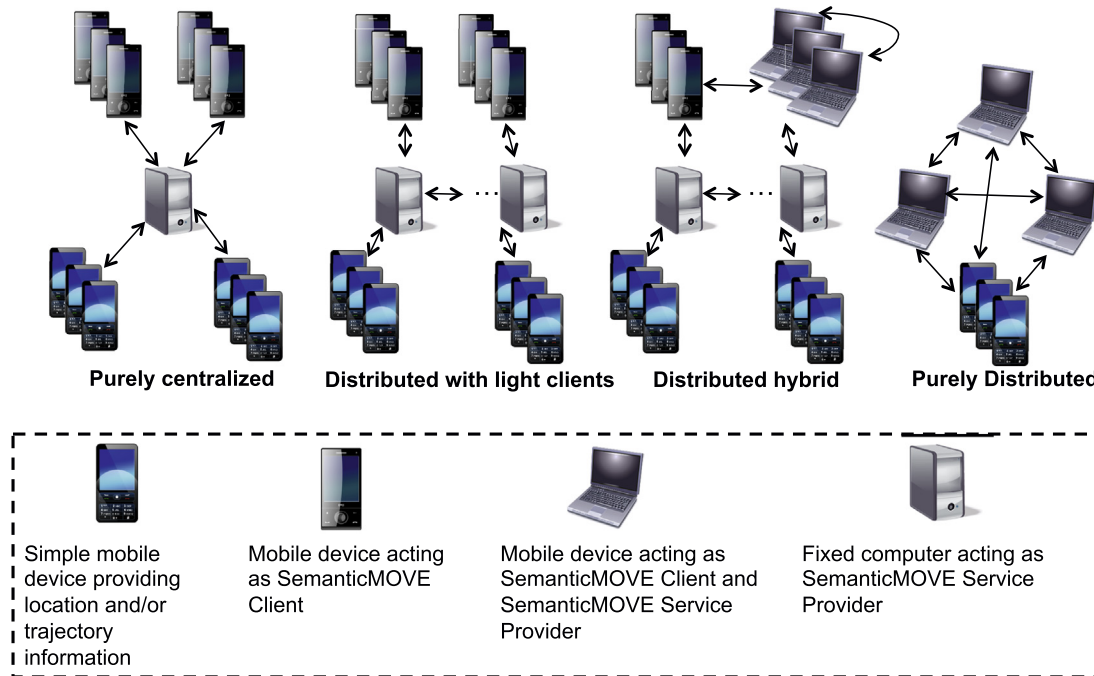
Fig. 4. Some possible deployments of SemanticMOVE.

information and assistance in various situations while preserving privacy. It should be noted that the proposed framework can accommodate small applications as well as generic and complex applications, although in the examples of Section 3 we mentioned the use of a number of specialized small applications.

5.3. Feasibility analysis

In Section 4, we analyzed research challenges that must be tackled to achieve a comprehensive management of moving objects and the semantics of mobility. In the two previous subsections, we have just described our framework proposal to address these challenges, as well as its deployment options. Now, we will provide some insights about the feasibility of the SemanticMOVE framework to support the development of real-world smart mobility systems and applications, summarized in Table 1.

We indicate the main challenges and difficulties to make the SemanticMOVE framework a reality. Besides, we briefly review some development approaches that can positively contribute and lead towards a full implementation of the SemanticMOVE framework providing its support for the development of smart mobility applications, such as those presented in Section 3. Some key research and development challenges that lie ahead are the following ones:

- *Collaborative data collection.* Participatory and mobile crowd-sensing are a hot research topic nowadays. The study presented in Ganti, Ye, and Lei (2011) indicates several research challenges in this area, such as: the pre-processing of data on the mobile devices (*localized analytics*) while coping with their limited resources (in terms of energy, bandwidth, and processing capabilities), the need to preserve the privacy of the participating users, the difficulty to ensure data integrity (e.g., to avoid the injection of false information), the difficulty of analyzing large amounts of data provided by a collection of mobile devices (*aggregate analytics*), and the interest of developing a unified common architecture rather than independent *application silos*. The study in Shilton (2012) indicates the challenges of

management, curation, and preservation of participatory personal data. The work presented in To, Ghinita, and Shahabi (2014) focuses on the problem of location privacy in *spatial crowdsourcing* scenarios (people carrying mobile devices perform certain data collection tasks –e.g., taking pictures of an accident– by physically moving to specific places). The challenges related to the collection, processing, and analysis of potentially-large amounts of streaming data obtained from sensors attached to a mobile device, as well as virtual sensors that represent user interaction with mobile applications and services, are emphasized in Brouwers and Langendoen (2012); Jayaraman, Perera, Georgakopoulos, and Zaslavsky (2013, 2014); Zaslavsky, Jayaraman, and Krishnaswamy (2013).

- *Semantic representation of moving objects.* Although a significant amount of work has been invested in the semantic modeling of moving objects, and particularly in the modeling of semantic trajectories, this is still an active research and development area. There are proposals that focus on the semantic representation of some aspects of moving objects (e.g., their trajectory, their context, the services they can access, etc.), but a unified approach that takes all those elements into account is still missing. Moreover, as we highlighted in Section 4.2, we think that the biggest challenge is the management of semantic models in a distributed way, where different moving objects may have complementary views of the world that need to be reconciled and combined.

- *Semantic query processing and reasoning.* As mentioned in Section 4.3, interesting work has been developed on context-awareness, efficient query processing on mobile environments, reasoning, and semantic search. But again, all these elements should be taken jointly into account, and an appropriate abstraction layer that enables the exploitation of the functionalities should be provided.

- *Analysis and mining of semantic trajectories.* As far as we know, large-scale spatio-temporal analytics is a quite unexplored research area. According to what we described in Section 4.4, inferring higher-level semantics (such as situations affecting a group of moving objects) from a large set of individual

Table 1
Summary of difficulties for the realization of the SemanticMOVE framework.

Challenges	Some sample contributors	Difficulties
(1) *Collaborative data collection*: representation, semantic enrichment, and on-device analysis of potentially-large amounts of data	Semantic annotation and fusion of sensor data: Calbimonte and Yan et al. (2012) and Zafeiropoulos et al. (2008) Mobile participatory sensing: Predic et al. (2013) Mobile crowd-sensing frameworks: Sherchan et al. (2012) and Jayaraman et al. (2013, 2014)	*Medium*. There is a need to ensure the scalability in managing large amounts of sensor data on the mobile devices, as well as to progress on automatic semantic annotation techniques. Further research on (peer-to-peer) P2P exchanges is also needed
(2) *Semantic representation of moving objects*: integrated semantic representation of all the aspects of moving objects (trajectory, context, accessible services, etc.), modular and distributed management of semantic models	Semantic location granules: Bernad et al. (2013) Semantic trajectories: Parent et al. (2013) and Fileto et al. (2013) Ontology mapping: Choi et al. (2006) Context modeling: Bettini et al. (2010) Semantic context management: Hu et al. (2009) and Konstantinou et al. (2007)	*Medium/High*. Several proposals concern the semantic representation of different aspects of moving objects, but the distributed management of semantic models, where moving objects may exchange and reconcile knowledge, is a quite unexplored area
(3) *Semantic query processing and reasoning*: integrated approach that combines context-awareness, efficient query processing on mobile environments, reasoning, and semantic search	Keyword-based queries: Trillo et al. (2007) Context reasoning: Bettini et al. (2010), Bikakis et al. (2008) and Tiberghien et al. (2012) Semantic reasoning on mobile devices: Ruta, Scioscia, and Sciascio et al. (2012), Yus et al. (2013), Bobed et al. (2014), Motik et al. (2012), Patton and McGuinness, 2014 (in press), Ruta et al. (2012) and Kim et al. (2010)	*Medium*. Several proposals have been developed for each of those subproblems. These areas are still under active research
(4) *Analysis and mining of semantic trajectories*: mobile analytics and large-scale spatio-temporal analytics	Aggregation and clustering of trajectories: Dodge et al. (2012) Mobility data mining: Haghighi et al. (2013), Predic and Stojanovic (2012) and Stahl et al. (2012) Mobile and ubiquitous semantic analytics: Ye et al. (2012) and Ermilov et al. (2014)	*High*. Although there is work on spatio-temporal data mining, it usually does not consider very large data sets. Performing analytics in a distributed environment on mobile devices is also a real challenge nowadays
(5) *Privacy protection*: preventing privacy attacks using semantic techniques	Fine-grained cloaking and semantic mobility privacy preservation: Damiani et al. (2011), Kapadia et al. (2008), Ghosh et al. (2012), Lee et al. (2011) and Celdran et al. (2014)	*High*. A higher amount of data combined with semantically-enhanced information and other background knowledge implies a higher privacy risk. How to ensure privacy in this environment does not seem clear for the moment

trajectories deserves further research, as well as exploiting those semantically-enhanced data for more powerful analysis. The research path widens considerably if we consider the interest of performing analysis tasks distributively on mobile devices.

- *Privacy protection*. This is another key pending challenge and a major concern in mobile computing in general. As explained in Section 4.5, the semantic enrichment of data increases the chances of privacy attacks, but at the same time semantic techniques are needed in order to ensure that potential attacks are prevented. This is an important difficulty, as it would require considering the different types of data sources and background knowledge that could be correlated to try to infer some information that should be kept private.

So, whereas there are still open research areas to make it a reality, we consider that further research and development advances could contribute to the realization of the proposed framework. Some of the existing challenges are related to the consideration of a distributed environment with heterogeneous moving objects, that may use different semantic representations and need to exchange data in a peer-to-peer way.

Some recent initiatives that could be considered quite aligned with the research issues considered in the SemanticMOVE framework are *SHERLOCK*[8] (*System for Heterogeneous mobilE Requests by Leveraging Ontological and Contextual Knowledge*) (Yus, Mena, Ilarri, & Illarramendi, 2013; Yus, Mena, Ilarri, & Illarramendi, in press) and the *SEEK* (*SEmantic Enrichment of trajectory Knowledge discovery*) project.[9] SHERLOCK is a distributed system whose goal is to provide

LBSs based on the use of semantic techniques and mobile agents (Milojicic et al., 1999; Trillo, Ilarri, & Mena, 2007). The participating objects/devices can cooperate and exchange data and knowledge among them to relieve the user from having to know, represent and use such knowledge himself/herself. The system exploits abductive and deductive reasoning to infer information to answer user requests continuously. SEEK focuses on researching methods to extract knowledge from large amounts of mobility data. As an example of the contributions of SEEK, the *Baquara* ontology presented in Fileto, Krüger, Pelekis, Theodoridis, and Renso (2013) provides a conceptual framework for the semantic enrichment of mobility data using annotations based on *Linked Data* (Bizer, Heath, & Berners-Lee, 2009).

6. Conclusions

In this position paper, we have outlined some research challenges that lie in the path towards smart mobility and we have presented SemanticMOVE, a framework designed for the semantic management of moving objects. The proposed distributed, general, and scalable framework can bring significant benefits. Semantic locations and trajectories enable the development of advanced LBSs that should provide more intelligent, proactive and valuable services to users navigating in outdoor and indoor environments. Moreover, the processing, analysis and mining of semantic locations and trajectories provide insights into the semantics of movement and the recognition of user activities, behaviors, and future movements.

SemanticMOVE benefits from the increasing sensing, processing, interaction, communication and energetic capabilities of mobile devices in a scalable way. As opposed to other related work, it is a quite generic architecture, supporting a fully distributed

---

environment for the management of the semantics of moving objects. So, each moving object can collect, represent, and analyze its own semantic mobility data and trajectories, and reason over them locally. Moving objects exchange with each other semantic mobility data, trajectories, and semantic concepts/knowledge, exploiting ad hoc wireless networks. The framework includes fixed servers only as an additional optional element, as the ad hoc cooperation among moving objects plays the fundamental role, rather than the traditional client/server interactions with centralized servers. We have presented two use case scenarios that illustrate the usefulness of the framework and show useful tasks that existing proposals cannot flexibly accomplish. Both small applications and generic and complex applications would take advantage of the functionalities of such a framework. The framework can be useful in a variety of mobile computing scenarios, such as in the case of smart cities (Ilarri, Stojanovic, & Ray, 2014).

Although the need and foundations of semantic management of mobility data have been already set up, significant research and innovations in all the identified challenges are needed to make our vision a reality. Some recent initiatives try to tackle the problems mentioned. However, more efforts are needed to solve all the challenges. We hope that this vision paper will encourage future work in these areas.

Acknowledgments

References

Afyouni, I., Ray, C., & Claramunt, C. (2012). Spatial models for indoor and context-aware navigation systems: A survey. *Journal of Spatial Information Science, 4*, 85–123.

Alonso, G. (2002). Myths around web services. *IEEE Data Engineering Bulletin, 25*, 3–9.

Alonso, G., Casati, F., Kuno, H., & Machiraju, V. (2004). *Web services – Concepts, architectures and applications.* Springer.

Alvares, L. O., Bogorny, V., Kuijpers, B., de Macedo, J. A. F., Moelans, B., & Vaisman, A. (2007). A model for enriching trajectories with semantic geographical information. In *15th Annual ACM International Symposium on Advances in Geographic Information Systems (GIS)* (pp. 22:1–22:8). ACM.

Andrienko, G., Andrienko, N., & Heurich, M. (2011). An event-based conceptual model for context-aware movement analysis. *International Journal of Geographical Information Science, 25*, 1347–1370.

Andrienko, G., Andrienko, N., Hurter, C., Rinzivillo, S., & Wrobel, S. (2013). Scalable analysis of movement data for extracting and exploring significant places. *IEEE Transactions on Visualization and Computer Graphics, 19*, 1078–1094.

Andrienko, G., Andrienko, N., & Wrobel, S. (2007). Visual analytics tools for analysis of movement data. *SIGKDD Explorations Newsletter, 9*, 38–46.

Antoniou, G., Corcho, O., Aberer, K., Simperl, E., & Studer, R. (2012). Semantic data management (Dagstuhl Seminar 12171). Dagstuhl reports, 2 (pp. 39–65).

Ardle, G. M., Petit, M., Ray, C., & Claramunt, C. (2012). Recommendations based on region and spatial profiles. In *11th International Symposium on Web and Wireless Geographical Information Systems (W2GIS). Lecture Notes in Computer Science (LNCS)* (Vol. 7236, pp. 167–184). Springer.

Ashbrook, D., & Starner, T. (2003). Using GPS to learn significant locations and predict movement across multiple users. *Personal and Ubiquitous Computing, 7*, 275–286.

Baglioni, M., Macedo, J., Renso, C., & Wachowicz, M. (2008). An ontology-based approach for the semantic modelling and reasoning on trajectories. In *Second International Workshop on Semantic and Conceptual Issues in Geographic Information Systems (SeCoGIS). Lecture Notes in Computer Science (LNCS)* (Vol. 5232, pp. 344–353). Springer.

Baldauf, M., Dustdar, S., & Rosenberg, F. (2007). A survey on context-aware systems. *International Journal of Ad Hoc and Ubiquitous Computing, 2*, 263–277.

Bao, J., Kendall, E. F., McGuinness, D. L., Patel-Schneider, P. F., Ding, L., & Khandelwal, A. (2012). OWL 2 web ontology language quick reference guide (2n ed.). W3C recommendation. <http://www.w3.org/TR/owl2-overview/>.

Barbieri, D. F., Braga, D., Ceri, S., Valle, E. D., & Grossniklaus, M. (2010). Querying RDF streams with C-SPARQL. *SIGMOD Record, 39*, 20–26.

Bechhofer, S., van Harmelen, F., Hendler, J., Horrocks, I., McGuinness, D. L., Patel-Schneider, P. F., & Stein, L. A. (2004). OWL web ontology language reference. W3C recommendation. <http://www.w3.org/TR/owl-ref/>.

Ben, T., Qin, X., & Wang, N. (2014). A semantic-based indexing for indoor moving objects. *International Journal of Distributed Sensor Networks, 2014*, 1–12.

Beresford, A. R., & Stajano, F. (2003). Location privacy in pervasive computing. *IEEE Pervasive Computing, 2*, 46–55.

Bernad, J., Bobed, C., Mena, E., & Ilarri, S. (2013). A formalization for semantic location granules. *International Journal of Geographical Information Science, 27*, 1090–1108. <CEUR-WS.org>.

Berners-Lee, T., Hendler, J., & Lassila, O. (2001). The semantic web. *Scientific American, 284*, 29–37.

Bettini, C., Brdiczka, O., Henricksen, K., Indulska, J., Nicklas, D., Ranganathan, A., et al. (2010). A survey of context modelling and reasoning techniques. *Pervasive and Mobile Computing, 6*, 161–180.

Bikakis, A., Patkos, T., Antoniou, G., & Plexousakis, D. (2008). A survey of semantics-based approaches for context reasoning in ambient intelligence. In *Constructing ambient intelligence, AmI (European Conference on Ambient Intelligence) 2007 Workshops. Communications in Computer and Information Science* (Vol. 11, pp. 14–23). Springer.

Biljecki, F., Ledoux, H., & van Oosterom, P. (2013). Transportation mode-based segmentation and classification of movement trajectories. *International Journal of Geographical Information Science, 27*, 385–407.

Bizer, C., Boncz, P., Brodie, M. L., & Erling, O. (2012). The meaningful use of Big Data: Four perspectives – four challenges. *SIGMOD Record, 40*, 56–60.

Bizer, C., Heath, T., & Berners-Lee, T. (2009). Linked Data – The story so far. *International Journal on Semantic Web and Information Systems, 5*, 1–22.

Bobed, C., Bobillo, F., Yus, R., Esteban, G., & Mena, E. (2014). Android went semantic: Time for evaluation. In *Third International Workshop on OWL Reasoner Evaluation (ORE). CEUR Workshop Proceedings* (Vol. 1207, pp. 23–29).

Bogorny, V., Avancini, H., de Paula, B. C., Kuplich, C. R., & Alvares, L. O. (2011). Weka-STPM: A software architecture and prototype for semantic trajectory data mining and visualization. *Transactions in GIS, 15*, 227–248.

Bogorny, V., Heuser, C. A., & Alvares, L. O. (2010). A conceptual data model for trajectory data mining. In *Sixth International Conference on Geographic Information Science (GIScience). Lecture Notes in Computer Science (LNCS)* (Vol. 6292, pp. 1–15). Springer.

Bogorny, V., Kuijpers, B., & Alvares, L. O. (2009). ST-DMQL: A semantic trajectory data mining query language. *International Journal of Geographical Information Science, 23*, 1245–1276.

Bogorny, V., Renso, C., de Aquino, A. R., de Lucca Siqueira, F., & Alvares, L. O. (2014). CONSTAnT – A conceptual data model for semantic trajectories of moving objects. *Transactions in GIS, 18*, 66–88.

Brouwers, N., & Langendoen, K. (2012). Pogo, a middleware for mobile phone sensing. In *13th International Middleware Conference (Middleware)* (pp. 21–40). Springer.

Calbimonte, J.-P., Yan, Z., Jeung, H., Corcho, Ó., & Aberer, K. (2012). Deriving semantic sensor metadata from raw measurements. In *Fifth International Workshop on Semantic Sensor Networks (SSN), CEUR Workshop Proceedings* (Vol. 904, pp. 33–48), <CEUR-WS.org>.

Calbimonte, J.-P., Jeung, H., Corcho, Ó., & Aberer, K. (2012). Enabling query technologies for the semantic sensor web. *International Journal on Semantic Web and Information Systems, 8*, 43–63.

Calvanese, D., Giacomo, G. D., & Lenzerini, M. (2002). A framework for ontology integration. In *The Emerging Semantic Web, selected papers from the First Semantic Web Working Symposium (SWWS). Frontiers in artificial intelligence and applications* (pp. 201–214). IOS Press.

Camossi, E., Villa, P., & Mazzola, L. (2013). Semantic-based anomalous pattern discovery in moving object trajectories. *CoRR, abs/1305.1946*, 1–20.

Cao, X., Cong, G., & Jensen, C. S. (2010). Mining significant semantic locations from GPS data. *Proceedings of the VLDB Endowment, 3*, 1009–1020.

Cao, H., & Wolfson, O. (2005). Nonmaterialized motion information in transport networks. In *10th International Conference on Database Theory (ICDT)* (pp. 173–188). Springer.

Castelli, G., Mamei, M., Rosi, A., & Zambonelli, F. (2009). Extracting high-level information from location data: The W4 diary example. *Mobile Networks and Applications, 14*, 107–119.

Celdran, A. H., Clemente, F. J. G., Perez, M. G., & Perez, G. M. (2014). SeCoMan: A semantic-aware policy framework for developing privacy-preserving and context-aware smart applications. *IEEE Systems Journal*, http://dx.doi.org/10.1109/JSYST.2013.2297707 [available online 21 January 2014].

Choi, N., Song, I.-Y., & Han, H. (2006). A survey on ontology mapping. *SIGMOD Record, 35*, 34–41.

Damiani, M. L., Bertino, E., & Silvestri, C. (2010). The PROBE framework for the personalized cloaking of private locations. *Transactions on Data Privacy, 3*, 123–148.

Damiani, M. L., Silvestri, C., & Bertino, E. (2011). Fine-grained cloaking of sensitive positions in location-sharing applications. *IEEE Pervasive Computing, 10*, 64–72.

d'Aquin, M., Sabou, M., & Motta, E. (2006). Modularization: A key for the dynamic selection of relevant knowledge components. In *First International Workshop on Modular Ontologies (WoMO), CEUR Workshop Proceedings* (Vol. 232, pp. 1–14). <CEUR-WS.org>.

Datta, S., Bhaduri, K., Giannella, C., Kargupta, H., & Wolff, R. (2006). Distributed data mining in peer-to-peer networks. *IEEE Internet Computing, 10*, 18–26.

Dentler, K., Cornet, R., ten Teije, A., & de Keizer, N. (2011). Comparison of reasoners for large ontologies in the OWL 2 EL profile. *Semantic Web, 2*, 71–87.

Dias, J. A., Rodrigues, J. J., & Zhou, L. (2014). Cooperation advances on vehicular communications: A survey. *Vehicular Communications, 1*, 22–32.

Ding, Z., & Deng, K. (2011). Collecting and managing network-matched trajectories of moving objects in databases. In *22nd International Conference on Database and Expert Systems Applications (DEXA)* (pp. 270–279). Springer.

Ding, L., Kolari, P., Ding, Z., & Avancha, S. (2007). Using ontologies in the semantic web: A survey. In *Ontologies. Integrated Series in Information Systems* (Vol. 14, pp. 79–113). Springer.

Dobre, C., & Xhafa, F. (2014). Intelligent services for Big Data science. *Future Generation Computer Systems, 37*, 267–281.

Dodge, S., Laube, P., & Weibel, R. (2012). Movement similarity assessment using symbolic representation of trajectories. *International Journal of Geographical Information Science, 26*, 1563–1588.

Dodge, S., Weibel, R., & Lautenschütz, A.-K. (2008). Towards a taxonomy of movement patterns. *Information Visualization, 7*, 240–252.

Do, T. M. T., & Gatica-Perez, D. (2014). The places of our lives: Visiting patterns and automatic labeling from longitudinal smartphone data. *IEEE Transactions on Mobile Computing, 13*, 638–648.

Ermilov, T., Khalili, A., & Auer, S. (2014). Ubiquitous semantic applications: A systematic literature review. *International Journal on Semantic Web and Information Systems, 10*, 66–99.

Fernández, J. D., Arias, M., Martínez-Prieto, M. A., & Gutiérrez, C. (2013). Management of big semantic data. In *Big data computing* (pp. 131–167). CRC Press.

Fileto, R., Krüger, M., Pelekis, N., Theodoridis, Y., & Renso, C. (2013). Baquara: A holistic ontological framework for movement analysis using linked data. In *32nd International Conference on Conceptual Modeling (ER). Lecture Notes in Computer Science (LNCS)* (Vol. 8217, pp. 342–355). Springer.

Fleming, B. (2013). Sensors – A forecast [automotive electronics]. *IEEE Vehicular Technology Magazine, 8*, 4–12.

Ganti, R. K., Ye, F., & Lei, H. (2011). Mobile crowdsensing: Current state and future challenges. *IEEE Communications Magazine, 49*, 32–39.

Ghosh, D., Joshi, A., Finin, T., & Jagtap, P. (2012). Privacy control in smart phones using semantically rich reasoning and context modeling. In *Symposium on Security and Privacy Workshops (SPW)* (pp. 82–85). IEEE Computer Society.

Giannotti, F., Nanni, M., Pedreschi, D., Pinelli, F., Renso, C., Rinzivillo, S., et al. (2011). Unveiling the complexity of human mobility by querying and mining massive trajectory data. *The VLDB Journal, 20*, 695–719.

Giannotti, F., Nanni, M., Pinelli, F., & Pedreschi, D. (2007). Trajectory pattern mining. In *13th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)* (pp. 330–339). ACM.

Giordano, S. (2002). Mobile ad hoc networks. In *Handbook of wireless networks and mobile computing* (pp. 325–346). John Wiley & Sons.

Görlach, A., Heinemann, A., & Terpstra, W. W. (2005). Survey on location privacy in pervasive computing. In *Privacy, security and trust within the context of pervasive computing. The International Series in Engineering and Computer Science* (Vol. 780, pp. 23–34). Springer.

Gruber, T. (1993). A translation approach to portable ontology specifications. *Knowledge Acquisition, 5*, 199–220.

Güting, R. H., Behr, T., & Düntgen, C. (2010). SECONDO: A platform for moving objects database research and for publishing and integrating research implementations. *IEEE Data Engineering Bulletin, 33*, 56–63.

Güting, R. H., de Almeida, T., & Ding, Z. (2006). Modeling and querying moving objects in networks. *The VLDB Journal, 15*, 165–190.

Haghighi, P. D., Krishnaswamy, S., Zaslavsky, A., Gaber, M. M., Sinha, A., & Gillick, B. (2013). Open Mobile Miner: A toolkit for building situation-aware data mining applications. *Journal of Organizational Computing and Electronic Commerce, 23*, 224–248.

Hartenstein, H., & Laberteaux, K. P. (2008). A tutorial survey on vehicular ad hoc networks. *IEEE Communications Magazine, 46*, 164–171.

Hornsby, K., & Egenhofer, M. J. (2002). Modeling moving objects over multiple granularities. *Annals of Mathematics and Artificial Intelligence, 36*, 177–194.

Horrocks, I. (2008). Ontologies and the semantic web. *Communications of the ACM, 51*, 58–67.

Hu, D. H., Dong, F., & Wang, C.-L. (2009). A semantic context management framework on mobile device. In *International Conference on Embedded Software and Systems (ICESS)* (pp. 331–338). IEEE Computer Society.

Hu, Y., Janowicz, K., Carral, D., Scheider, S., Kuhn, W., Berg-Cross, G., et al. (2013). A geo-ontology design pattern for semantic trajectories. In *11th International Conference on Spatial Information Theory (COSIT). Lecture Notes in Computer Science (LNCS)* (Vol. 8116, pp. 438–456). Springer.

Ilarri, S., Bobed, C., & Mena, E. (2011). An approach to process continuous location-dependent queries on moving objects with support for location granules. *Journal of Systems and Software, 84*, 1327–1350.

Ilarri, S., Stojanovic, D., & Ray, C. (2014). Semantic management of moving objects in smart cities. ERCIM News, Special theme on smart cities (pp. 22–23), 2014.

Ilarri, S., Illarramendi, A., Mena, E., & Sheth, A. (2011). Semantics in Location-Based Services – Guest editors' introduction for special issue. *IEEE Internet Computing, 15*, 10–14.

Ilarri, S., Mena, E., & Illarramendi, A. (2006). Location-dependent queries in mobile contexts: Distributed processing using mobile agents. *IEEE Transactions on Mobile Computing, 5*, 1029–1043.

Ilarri, S., Mena, E., & Illarramendi, A. (2010). Location-dependent query processing: Where we are and where we are heading. *ACM Computing Surveys, 42*, 12:1–12:73.

Jayaraman, P. P., Perera, C., Georgakopoulos, D., & Zaslavsky, A. (2013). MOSDEN: An internet of things middleware for resource constrained mobile devices. *CoRR, abs/1310.4038*, 1–10.

Jayaraman, P. P., Perera, C., Georgakopoulos, D., & Zaslavsky, A. (2014). MOSDEN: A scalable mobile collaborative platform for opportunistic sensing applications. *EAI Endorsed Transactions on Collaborative Computing, 1*, 1–16.

Jiménez-Ruiz, E., Grau, B. C., Sattler, U., Schneider, T., & Berlanga, R. (2008). Safe and economic re-use of ontologies: A logic-based methodology and tool support. In *Fifth European Semantic Web Conference (ESWC). Lecture Notes in Computer Science (LNCS)* (Vol. 5021, pp. 185–199). Springer.

Kang, J. H., Welbourne, W., Stewart, B., & Borriello, G. (2004). Extracting places from traces of locations. In *Second ACM International Workshop on Wireless Mobile Applications and Services on WLAN Hotspots (WMASH)* (pp. 110–118). ACM.

Kapadia, A., Triandopoulos, N., Cornelius, C., Peebles, D., & Kotz, D. (2008). AnonySense: Opportunistic and privacy-preserving context collection. In *Sixth International Conference on Pervasive Computing (Pervasive). Lecture Notes in Computer Science (LNCS)* (Vol. 5013, pp. 280–297). Springer.

Kim, T., Park, I., Hyun, S., & Lee, D. (2010). MiRE4OWL: Mobile rule engine for OWL. In *34th Computer Software and Applications Conference Workshops (COMPSACW)* (pp. 317–322). IEEE Computer Society.

Konstantinou, N., Solidakis, E., Zoi, S., Zafeiropoulos, A., Stathopoulos, P., & Mitrou, N. (2007). Priamos: A middleware architecture for real-time semantic annotation of context features. In *Third IET International Conference on Intelligent Environments (IE)* (pp. 96–103). IEEE Computer Society.

Krumm, J. (2009). A survey of computational location privacy. *Personal and Ubiquitous Computing, 13*, 391–399.

Krumm, J., & Horvitz, E. (2006). Predestination: Inferring destinations from partial trajectories. In *Eighth International Conference on Ubiquitous Computing (UbiComp)* (pp. 243–260). Springer.

Kumar, A., Kantardzic, M., & Madden, S. (2006). Guest editors' introduction: Distributed data mining – framework and implementations. *IEEE Internet Computing, 10*, 15–17.

Lane, N. D. (2012). Community-aware smartphone sensing systems. *IEEE Internet Computing, 16*, 60–64.

Lane, N. D., Miluzzo, E., Lu, H., Peebles, D., Choudhury, T., & Campbell, A. T. (2010). A survey of mobile phone sensing. *IEEE Communications Magazine, 48*, 140–150.

Laube, P. (2009). Progress in movement pattern analysis. In *Behaviour monitoring and interpretation – BMI, Ambient intelligence and Smart Environments* (Vol. 3, pp. 43–71). IOS Press.

Laube, P., Imfeld, S., & Weibel, R. (2005). Discovering relative motion patterns in groups of moving point objects. *International Journal of Geographical Information Science, 19*, 639–668.

Lee, J.-G., Han, J., Li, X., & Gonzalez, H. (2008). TraClass: Trajectory classification using hierarchical region-based and trajectory-based clustering. *Proceedings of the VLDB Endowment, 1*, 1081–1094.

Lee, J.-G., Han, J., & Whang, K.-Y. (2007). Trajectory clustering: A partition-and-group framework. In *International Conference on Management of Data (SIGMOD)* (pp. 593–604). ACM.

Lee, B., Oh, J., Yu, H., & Kim, J. (2011). Protecting location privacy using location semantics. In *17th International Conference on Knowledge Discovery and Data Mining (SIGKDD)* (pp. 1289–1297). ACM.

Liao, L., Fox, D., & Kautz, H. (2005). Location-based activity recognition using relational markov networks. In *19th International Joint Conference on Artificial Intelligence (IJCAI)* (pp. 773–778). Morgan Kaufman.

Li, X., Claramunt, C., Ray, C., & Lin, H. (2006). A semantic-based approach to the representation of network-constrained trajectory data. In *12th International Symposium on Spatial Data Handling (SDH)* (pp. 451–464). Springer.

Li, Z., Han, J., Ji, M., Tang, L. A., Yu, Y., Ding, B., et al. (2011). MoveMine: Mining moving object data for discovery of animal movement patterns. *ACM Transactions on Intelligent Systems and Technology, 2*, 37:1–37:32.

Lim, L., Wang, H., & Wang, M. (2007). Semantic data management: Towards querying data with their meaning. In *23rd International Conference on Data Engineering (ICDE)* (pp. 1438–1442). IEEE Computer Society.

Liu, N., Liu, M., Cao, J., Chen, G., & Lou, W. (2010). When transportation meets communication: V2P over VANETs. In *30th International Conference on Distributed Computing Systems (ICDCS)* (pp. 567–576). IEEE Computer Society.

Mandreoli, F., Martoglia, R., Penzo, W., & Sassatelli, S. (2009). Data-sharing P2P networks with semantic approximation capabilities. *IEEE Internet Computing, 13*, 60–70.

McIlraith, S. A., Son, T. C., & Zeng, H. (2001). Semantic web services. *IEEE Intelligent Systems, 16*, 46–53.

Milojicic, D., Douglis, F., & Wheeler, R. (Eds.). (1999). *Mobility: Processes, computers, and agents*. ACM Press/Addison-Wesley Publishing Co.

Mishra, R. B., & Kumar, S. (2011). Semantic web reasoners and languages. *Artificial Intelligence Review, 35*, 339–368.

Monreale, A., Pinelli, F., Trasarti, R., & Giannotti, F. (2009). WhereNext: A location predictor on trajectory pattern mining. In *15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)* (pp. 637–646). ACM.

Monreale, A., Trasarti, R., Pedreschi, D., Renso, C., & Bogorny, V. (2011). C-safety: A framework for the anonymization of semantic trajectories. *Transactions on Data Privacy, 4*, 73–101.

Motik, B., Horrocks, I., & Kim, S. M. (2012). Delta-reasoner: A semantic web reasoner for an intelligent mobile platform. In *21st International World Wide Web Conference (WWW)* (pp. 63–72). ACM.

Mountain, D., & Raper, J. (2001). Modelling human spatio-temporal behaviour: A challenge for location-based services. In *Sixth International Conference on Geocomputation (Geocomputation)*, 2001, 9 pages.

Musolesi, M. (2014). Big mobile data mining: Good or evil? *IEEE Internet Computing, 18*, 78–81.

Nanni, M., Kuijpers, B., Koerner, C., May, M., & Pedreschi, D. (2008). Spatiotemporal data mining. In *Mobility, data mining and privacy* (pp. 267–296). Springer.

Navigli, R. (2009). Word sense disambiguation: A survey. *ACM Computing Surveys, 41*, 10:1–10:69.

Olariu, S., & Weigle, M. C. (2009). *Vehicular networks: From theory to practice*. CRC Press.

Palma, A. T., Bogorny, V., Kuijpers, B., & Alvares, L. O. (2008). A clustering-based approach for discovering interesting places in trajectories. In *Symposium on Applied Computing (SAC)* (pp. 863–868). ACM.

Parent, C., Spaccapietra, S., Renso, C., Andrienko, G., Andrienko, N., Bogorny, V., et al. (2013). Semantic trajectories modeling and analysis. *ACM Computing Surveys, 45*, 42:1–42:32.

Park, J.-H. (2014). Spatial semantic search in location-based web services. In *23rd International World Wide Web Conference (WWW)* (pp. 9–14). ACM.

Patton, E., & McGuinness, D. (in press). A power consumption benchmark for reasoners on mobile devices. In *The 13th International Semantic Web Conference (ISWC)*, 2014.

Payne, T. R., & Lassila, O. (2004). Guest editors' introduction: Semantic web services. *IEEE Intelligent Systems, 19*, 14–15.

Pejovic, V., & Musolesi, M. (2013). Anticipatory mobile computing: A survey of the state of the art and research challenges. Technical Report CSR-13-02, School of Computer Science, University of Birmingham.

Pelekis, N., Andrienko, G., Andrienko, N., Kopanakis, I., Marketos, G., & Theodoridis, Y. (2012). Visually exploring movement data via similarity-based analysis. *Journal of Intelligent Information Systems, 38*, 343–391.

Pelekis, N., Frentzos, E., Giatrakos, N., & Theodoridis, Y. (2008). HERMES: Aggregative LBS via a trajectory DB engine. In *International Conference on Management of Data (SIGMOD)* (pp. 1255–1258). ACM.

Pelekis, N., Theodoridis, Y., & Janssens, D. (2014). On the management and analysis of our lifesteps. *ACM SIGKDD Explorations Newsletter, 15*, 23–32.

Perera, C., Zaslavsky, A. B., Christen, P., & Georgakopoulos, D. (2014). Context aware computing for the internet of things: A survey. *IEEE Communications Surveys & Tutorials, 16*, 414–454.

Predic, B., & Stojanovic, D. (2012). Localized processing and analysis of accelerometer data in detecting traffic events and driver behaviour. *Journal of Universal Computer Science, 18*, 1152–1176.

Predic, B., Yan, Z., Eberle, J., Stojanovic, D., & Aberer, K. (2013). ExposureSense: Integrating daily activities with air quality using mobile participatory sensing. In *11th IEEE Pervasive Computing and Communication (PerCom) Workshops* (pp. 303–305). IEEE Computer Society.

Reddy, S., Samanta, V., Burke, J., Estrin, D., Hansen, M. H., & Srivastava, M. B. (2009). MobiSense – Mobile network services for coordinated participatory sensing. In *International Symposium on Autonomous Decentralized Systems (ISADS)* (pp. 231–236). IEEE Computer Society.

Renso, C., Baglioni, M., de Macedo, J. A. F., Trasarti, R., & MonicaWachowicz (2013). How you move reveals who you are: Understanding human behavior by analyzing trajectory data. *Knowledge and Information Systems, 37*, 331–362.

Richter, K.-F., Schmid, F., & Laube, P. (2012). Semantic trajectory compression: Representing urban movement in a nutshell. *Journal of Spatial Information Science*, 3–30.

Rocha, J. A. M. R., Oliveira, G., Alvares, L. O., Bogorny, V., & Times, V. C. (2010). DB-SMoT: A direction-based spatio-temporal clustering method. In *Fifth IEEE International Conference on Intelligent Systems (IS)* (pp. 114–119). IEEE Computer Society.

Ruta, M., Scioscia, F., Sciascio, E. D., Gramegna, F., & Loseto, G. (2012). Mini-ME: The mini matchmaking engine. In *First International Workshop on OWL Reasoner Evaluation (ORE)*, CEUR Workshop Proceedings (Vol. 858, pp. 1–12), <CEUR-WS.org>.

Ruta, M., Scioscia, F., Loseto, G., Gramegna, F., & Sciascio, E. D. (2012). A mobile reasoner for semantic-based matchmaking. In *Sixth International Conference on Web Reasoning and Rule Systems (RR). Lecture Notes in Computer Science (LNCS)* (Vol. 7497, pp. 254–257). Springer.

Schiller, J., & Voisard, A. (2004). *Location-based services*. Elsevier.

Shadbolt, N., Berners-Lee, T., & Hall, W. (2006). The Semantic Web revisited. *IEEE Intelligent Systems, 21*, 96–101.

Shankar, P., Huang, Y.-W., Castro, P., Nath, B., & Iftode, L. (2012). Crowds replace experts: Building better location-based services using mobile social network interactions. In *Tenth IEEE International Conference on Pervasive Computing and Communications (PerCom)* (pp. 20–29). IEEE Computer Society.

Sherchan, W., Jayaraman, P. P., Krishnaswamy, S., Zaslavsky, A., Loke, S., & Sinha, A. (2012). Using on-the-move mining for mobile crowdsensing. In *13th International Conference on Mobile Data Management (MDM)* (pp. 115–124). IEEE Computer Society.

Sheth, A., Henson, C., & Sahoo, S. S. (2008). Semantic sensor web. *IEEE Internet Computing, 12*, 78–83.

Sheth, A., & Perry, M. (2008). Traveling the semantic web through space, time, and theme. *IEEE Internet Computing, 12*, 81–86.

Shilton, K. (2012). Participatory personal data: An emerging research challenge for the information sciences. *Journal of the American Society for Information Science and Technology, 63*, 1905–1915.

Spaccapietra, S., & Parent, C. (2011). Adding meaning to your steps. In *30th International Conference on Conceptual Modeling (ER)* (pp. 13–31). Springer.

Spaccapietra, S., Parent, C., Damiani, M. L., de Macedo, J. A., Porto, F., & Vangenot, C. (2008). A conceptual view on trajectories. *Data & Knowledge Engineering, 65*, 126–146.

Spinsant, L., Celli, F., & Renso, C. (2010). Where you stop is who you are: Understanding people's activities by places visited. In *Fifth Workshop on Behaviour Monitoring and Interpretation (BMI)*.

Stahl, F., Gaber, M. M., Aldridge, P., May, D., Liu, H., Bramer, M., et al. (2012). Homogeneous and heterogeneous distributed classification for pocket data mining. *Transactions on Large-Scale Data- and Knowledge-Centered Systems, 7100*, 183–205.

Sui, D., Elwood, S., & Goodchild, M. (2012). *Crowdsourcing geographic knowledge: volunteered geographic information (VGI) in theory and practice*. Springer.

Su, H., Zheng, K., Zheng, K., Huang, J., & Zhou, X. (2014). STMaker – a system to make sense of trajectory data. *Proceedings of the VLDB Endowment, 7*, 1701–1704.

Tiberghien, T., Mokhtari, M., Aloulou, H., & Biswas, J. (2012). Semantic reasoning in context-aware assistive environments to support ageing with dementia. In *11th International Semantic Web Conference (ISWC). Lecture Notes in Computer Science (LNCS)* (Vol. 7650, pp. 212–227). Springer.

To, H., Ghinita, G., & Shahabi, C. (2014). A framework for protecting worker location privacy in spatial crowdsourcing. In *40th International Conference on Very Large Data Bases (VLDB)* (pp. 919–930). VLDB Endowment.

Trillo, R., Gracia, J., Espinoza, M., & Mena, E. (2007). Discovering the semantics of user keywords. *Journal of Universal Computer Science, 13*, 1908–1935.

Trillo, R., Ilarri, S., & Mena, E. (2007). Comparison and performance evaluation of mobile agent platforms. In *Third International Conference on Autonomic and Autonomous Systems (ICAS)* (pp. 41–46). IEEE Computer Society.

Tsoumakas, G., & Vlahavas, I. (2009). Distributed data mining. In *Encyclopedia of data warehousing and mining* (pp. 709–715). IGI Global.

Uddin, M. R., Ravishankar, C., & Tsotras, V. J. (2011). Finding regions of interest from trajectory data. In *12th International Conference on Mobile Data Management (MDM)* (pp. 39–48). IEEE Computer Society.

Uschold, M., & Gruninger, M. (2004). Ontologies and semantics for seamless connectivity. *SIGMOD Record, 33*, 58–64.

Valdés, F., Damiani, M. L., & Güting, R. H. (2013). Symbolic trajectories in SECONDO: Pattern matching and rewriting. In *18th International Conference on Database Systems for Advanced Applications (DASFAA). Lecture Notes in Computer Science (LNCS)* (Vol. 7826, pp. 450–453). Springer.

Valle, E. D., Ceri, S., van Harmelen, F., & Fensel, D. (2009). It's a streaming world! Reasoning upon rapidly changing information. *IEEE Intelligent Systems, 24*, 83–89.

Vazirgiannis, M., & Wolfson, O. (2001). A spatiotemporal model and language for moving objects on road networks. In *Seventh International Symposium on Advances in Spatial and Temporal Databases (SSTD). Lecture Notes in Computer Science (LNCS)* (Vol. 2121, pp. 20–35). Springer.

Vicente, C. R., Freni, D., Bettini, C., & Jensen, C. S. (2011). Location-related privacy in geo-social networks. *IEEE Internet Computing, 15*, 20–27.

Viktoratos, I., Tsadiras, A. K., & Bassiliades, N. (2014). Using rules to develop a personalized and social location information system for the semantic web. In *Eighth International Web Rule Symposium (RuleML). Lecture Notes in Computer Science (LNCS)* (Vol. 8620, pp. 82–96). Springer.

Wang, X., Li, G., Jiang, G., & Shi, Z. (2013). Semantic trajectory-based event detection and event pattern mining. *Knowledge and Information Systems, 37*, 305–329.

Weiser, M. (1991). The computer for the 21st century. *Scientific American, 265*, 94–104.

Wirz, M., Mitleton-Kelly, E., Franke, T., Camilleri, V., Montebello, M., Roggen, D., et al. (2013). Using mobile technology and a participatory sensing approach for crowd monitoring and management during large-scale mass gatherings. In *Co-evolution of intelligent socio-technical systems understanding complex systems* (pp. 61–77). Springer.

Wolfson, O., & Mena, E. (2004). Applications of moving objects databases. In *Spatial databases: Technologies, techniques and trends* (pp. 186–203). Idea Group, Inc.

Wolfson, O., Xu, B., Chamberlain, S., & Jiang, L. (1998). Moving objects databases: Issues and solutions. In *Tenth International Conference on Scientific and Statistical Database Management (SSDM)* (pp. 111–122). IEEE Computer Society.

Wolfson, O., Xu, B., & Sistla, A. P. (2004). An economic model for resource exchange in mobile peer to peer networks. In *16th International Conference on Scientific and Statistical Database Management (SSDM)* (pp. 235–244). IEEE Computer Society.

Wood, Z., & Galton, A. (2009). Classifying collective motion. In *Behaviour monitoring and interpretation – BMI, Ambient Intelligence and Smart Environments* (Vol. 3, pp. 129–155). IOS Press.

Wu, F., Lei, T. K. H., Li, Z., & Han, J. (2014). MoveMine 2.0: Mining object relationships from movement data. *Proceedings of the VLDB Endowment, 7*, 1613–1616.

Xie, K., Deng, K., & Zhou, X. (2009). From trajectories to activities: A spatio-temporal join approach. In *International Workshop on Location Based Social Networks (LBSN)* (pp. 25–32). ACM.

Xu, J., & Güting, R. H. (2013). A generic data model for moving objects. *Geoinformatica, 17*, 125–172.

Xu, B., Wolfson, O., & Rishe, N. (2006). Benefit and pricing of spatio-temporal information in mobile peer-to-peer networks. In *39th Annual Hawaii International Conference on System Sciences (HICSS)* (pp. 1–10). IEEE Computer Society.

Yan, Z., Chakraborty, D., Parent, C., Spaccapietra, S., & Aberer, K. (2011). SeMiTri: A framework for semantic annotation of heterogeneous trajectories. In *14th International Conference on Extending Database Technology (EDBT)* (pp. 259–270). ACM.

Yan, Z., Giatrakos, N., Katsikaros, V., Pelekis, N., & Theodoridis, Y. (2011). SeTraStream: Semantic-aware trajectory construction over streaming movement data. In *12th International Conference on Advances in Spatial and Temporal Databases (SSTD). Lecture Notes in Computer Science (LNCS)* (Vol. 6849, pp. 367–385). Springer.

Yan, Z., Macedo, J., Parent, C., & Spaccapietra, S. (2008). Trajectory ontologies and queries. *Transactions in GIS, 12*, 75–91.

Yan, Z., Parent, C., Spaccapietra, S., & Chakraborty, D. (2010). A hybrid model and computing platform for spatio-semantic trajectories. In *Seventh Extended Semantic Web Conference (ESWC). Lecture Notes in Computer Science (LNCS)* (Vol. 6088, pp. 60–75). Springer.

Ye, J., Dobson, S., & McKeever, S. (2012). Situation identification techniques in pervasive computing: A review. *Pervasive and Mobile Computing, 8*, 36–66.

Yigitoglu, E., Damiani, M. L., Abul, O., & Silvestri, C. (2012). Privacy-preserving sharing of sensitive semantic locations under road-network constraints. In *13th International Conference on Mobile Data Management (MDM)* (pp. 186–195). IEEE Computer Society.

yi Hong, J., ho Suh, E., & Kim, S.-J. (2009). Context-aware systems: A literature review and classification. *Expert Systems with Applications, 36*, 8509–8522.

Ying, J.-C., Chen, H.-S., Lin, K. W., Lu, E. H.-C., Tseng, V. S., Tsai, H.-W., et al. (2014). Semantic trajectory-based high utility item recommendation system. *Expert Systems with Applications, 41*, 4762–4776.

Yus, R., Bobed, C., Esteban, G., Bobillo, F., & Mena, E., (2013). Android goes semantic: DL reasoners on smartphones. In *Second International Workshop on OWL Reasoner Evaluation, CEUR Workshop Proceedings (ORE)* (Vol. 1015, pp. 46–52), 2013. <CEUR-WS.org>.

Yus, R., Mena, E., Ilarri, S., & Illarramendi, A. (in press). SHERLOCK: Semantic management of location-based services in wireless environments. *Pervasive and Mobile Computing.* http://dx.doi.org/10.1016/j.pmcj.2013.07.018 [corrected proof, available online 13 August 2013].

Yus, R., Mena, E., Ilarri, S., & Illarramendi, A. (2013). SHERLOCK: A system for location-based services in wireless environments using semantics. In *22nd International World Wide Web Conference (WWW)* (pp. 301–304). ACM.

Yu, S., & Spaccapietra, S. (2010). A knowledge infrastructure for intelligent query answering in location-based services. *GeoInformatica, 14*, 379–404.

Yu, M.-C., Yu, T., Wang, S.-C., Lin, C.-J., & Chang, E. Y. (2014). Big data small footprint: The design of a low-power classifier for detecting transportation modes. *Proceedings of the VLDB Endowment, 7*, 1429–1440.

Zafeiropoulos, A., Konstantinou, N., Arkoulis, S., Spanos, D.-E., & Mitrou, N. (2008). A semantic-based architecture for sensor data fusion. In *Second International Conference on Mobile Ubiquitous Computing, Systems, Services and Technologies (UBICOMM)* (pp. 116–121). IEEE Computer Society.

Zaslavsky, A., Jayaraman, P., & Krishnaswamy, S. (2013). ShareLikesCrowd: Mobile analytics for participatory sensing and crowd-sourcing applications. In *29th International Conference on Data Engineering Workshops (ICDEW)* (pp. 128–135). IEEE Computer Society.

Zaslavsky, A. B., Perera, C., & Georgakopoulos, D. (2013). Sensing as a service and Big Data. *CoRR, abs/1301.0159*, 1–8.

Zheng, Y., Zhang, L., Ma, Z., Xie, X., & Ma, W.-Y. (2011). Recommending friends and locations based on individual location history. *ACM Transactions on the Web, 5*, 5:1–5:44.