Automatically Tailoring Semantics-enabled Dimensions for Movement Data Warehouses

Juarez A. P. Sacenti[1], Fabio Salvini[2], Renato Fileto[1], Alessandra Raffaetà[2], and Alessandro Roncato[2]

[1] PPGCC/INE/CTC, Federal University of Santa Catarina, Florianópolis-SC, Brazil
juarez.sacenti@posgrad.ufsc.br, r.fileto@ufsc.br
[2] DAIS, Università Ca' Foscari Venezia, Venezia, Italy
salvini.fabio001@gmail.com, raffaeta@unive.it, roncato@unive.it

Abstract. This paper proposes an automatic approach to build tailored dimensions for movement data warehouses based on views of existing hierarchies of objects (and their respective classes) used to semantically annotate movement segments. It selects the objects (classes) that annotate at least a given number of segments of a movement dataset to delineate hierarchy views for deriving tailored analysis dimensions for that movement dataset. Dimensions produced in this way can be quite smaller than the hierarchies from which they are extracted, leading to efficiency gains, among other potential benefits. Results of experiments with tweets semantically enriched with points of interest taken from linked open data collections show the viability of the proposed approach.
The final authenticated version is available online at http://dx.doi.org/10.1007/978-3-319-22729-0_16.

Keywords: movement data, data warehouses, analysis dimensions, semantic web, Linked Open Data (LOD), social media

1 Introduction

Nowadays, it is possible to accumulate and exploit large volumes of movement data, such as moving objects' trajectories gathered by using positioning technologies (e.g. GPS), and sequences of users' posts on social media [1, 2]. On top of this, new methods have been developed to semantically enrich such data, for example with objects (instances) and concepts (classes of objects) that help to describe what goes on with the moving objects (e.g., places visited, transportation means employed, activities performed, goals of stops and moves) [3–6]. These developments unleash opportunities to build Movement Data Warehouses (MDWs), i.e., DWs of semantically enriched movement data [7–9], for a myriad of application fields, ranging from traffic and homeland security to geographically concentrated marketing and social behavior studies. However, to realize MDWs, besides accommodating large volumes of semantically annotated spatio-temporal data in the dimensional model [9], it is still necessary to develop appropriate ways to exploit semantic annotations for semantics-enabled movement data analysis.

This paper makes contributions towards building MDWs that support movement data analysis founded on their semantic annotations. It introduces an automatic approach to build tailored analysis dimensions from hierarchies of objects or classes used to semantically annotate movement segments, i.e., specific positions or sequences of positions occupied by moving objects, such as stops and moves [3]. For instance, a move can be semantically annotated with a concept that describes the transportation means employed (e.g., taxi, bus, tram) in a conceptual hierarchy or taxonomy. A stop can be semantically associated via some property (e.g. visits) with a Place of Interest (PoI). The PoIs are geographic objects that can be hierarchically organized according to their *part_of* relationships (such as the containment of a *shop* in a *shopping mall* that is in a particular *neighborhood* of a *city*). Notice that classes of objects such as PoIs can be regarded as concepts or categories (e.g., *city*, *neighborhood*, *shopping mall*, *shop*, *clothes shop*, *pet shop*, *barber shop*) and organized in another hierarchy determined by *is_a* relationships. Both, hierarchies of objects and hierarchies of their respective classes can be useful for information analysis in MDWs [9].

In this work, we consider semantic annotations of movement segments as associations with objects and concepts of Linked Open Data (LOD). Such associations are produced by methods such as those proposed in [4] and [6]. Hierarchies of LOD objects and hierarchies of their respective classes can be obtained by exploiting chains of partially ordering relationships, such as *part_of* relationships between objects and *is_a* relationships between concepts. They can also be built or complemented by relationships that are inferred from existing ones or data properties (e.g., geographic extensions). Then, their objects (classes) that annotate at least a given number of segments of a Movement Dataset (MoD) are selected and the hierarchy is traversed from them up to their least common ancestor, to produce tree-like hierarchy views which are the basis to build tailored analysis dimensions for that MoD.

Results of experiments with tweets taken via the Twitter API[3] and semantically enriched with PoIs of the DBpedia[4] and LinkedGeoData (LGD)[5] LOD collections testify the viability and some benefits of the proposed approach. The dimensions produced by our technique can be quite smaller than the hierarchies present in LOD from which they are extracted. Although we demonstrate our proposal on social media trails semantically enriched with PoIs of LOD collections, the proposed approach could be applied to other kinds of movement data enriched with other collections of objects and/or concepts, as it relies just on the number of hits of each object and concept, i.e., the number of times the object/concept is used to annotate segments of a particular MoD.

The rest of this paper is structured as follows. Section 2 provides basic definitions. Section 3 describes how to extract hierarchies from LOD collections, and how to tailor such hierarchies into MDWs' dimensions. Section 4 reports experiments that apply the proposed approach to build spatial dimensions of

[3] https://dev.twitter.com/rest/public
[4] http://dbpedia.org
[5] http://linkedgeodata.org

PoIs and their categories taken from the LOD collections used to semantically annotate tweets. Section 5 discusses related work. Finally, Section 6 concludes the paper, by pointing out its contributions and enumerating future work.

2 Basic definitions

This section introduces the definitions needed to understand our method for tailoring dimensions for MDWs. They describe the inputs of the proposed method, namely: movement data segments, semantically enriching resources (objects and their classes), associations of movement data with these resources to annotate the former, and resources hierarchies. A **Movement Dataset (MoD)** is a set of movement data segments like the one formally described in Definition 1.

Definition 1. *A* ***Movement Data Segment*** *(MDS) is a time-ordered sequence* $\langle p_1, \ldots, p_n \rangle$ *(*$n \in \mathbb{N}^+$*) of spatio-temporal positions of a moving object* MO*, where each position is a tuple* $p_i = (x_i, y_i, t_i)$ *with* $1 \leq i \leq n$*, where* (x_i, y_i) *are geographic coordinates, and* t_i *is a timestamp indicating the instant when* MO *was at those coordinates.*

An MDS can be any sequence of spatio-temporal coordinates occupied by a moving object, such as trajectory positions, geo-located user's posts on a social media, or geo-referenced logs of any kind of information system (e.g., Web logs). An MDS can be segmented into smaller MDSs at several abstraction levels (e.g., first in sub-trajectories that can refer to specific trips, and then in episodes such as stops and moves [3]). The resulting movement segments can be hierarchically organized according to containment relationships of their time spans, for information analysis in DWs as proposed in [9]. The movement segments can be enriched with semantic annotations to describe them according to several dimensions, such as those proposed in [5], namely space, time, characteristics of the moving objects, activities performed by these objects, their goal, transportation means employed, and the relevant environmental conditions during each movement segment time span. Such a semantic annotation can be an association with a resource (object or class) as stated by Definition 2.

Definition 2. *A* ***Semantically Enriching Resource*** *is a concept or instance of concept that can be used to semantically annotate an MDS.*

A semantically enriching resource (or simply resource) can be implemented as a reference (e.g., a Universal Resource Identifier - URI) to a class (concept) or object (instance) defined in an ontology, knowledge base, or even a LOD collection. In the latter a resource is denoted by a URI and relationships between resources are expressed by RDF triples.

For example, in Fig. 1 on the right some LOD resources are shown: the classes *restaurant* and *city*, two objects of the class restaurant, labeled *Ae Oche* and *Sushi Wok*, and an object of the class *city* named *Mestre*. These resources are used to annotate the tweets on the left side of the figure which are spatio-temporal positions. In fact they contain spatial coordinates expressed by longitude and latitude (in the lower right corner) and a temporal information specifying where and when tweets have been sent. The property that links tweets to resources is *visits*, to indicate that the tweets have been sent from the places (resources) where the person is.

LOD resources are convenient to semantically annotate movement data, as they have well-defined semantics and can be integrated more easily thanks to their adherence to semantic Web standards. Many information collections are publicly available in the Web by following Linked Data principles [10] (e.g., DBpedia, LGD). In addition, they have been continuously growing and kept up to date by active Web communities, which link their resources via a variety of properties, that include *owl:sameAs* (for objects duplicated in distinct collections), and *rdfs:equivalentClass* (for duplicated concepts). A direct association annotates an MDS with a property (e.g., visits) whose property value is a resource (e.g., the *restaurant Ae Oche*), as stated by Definition 3.

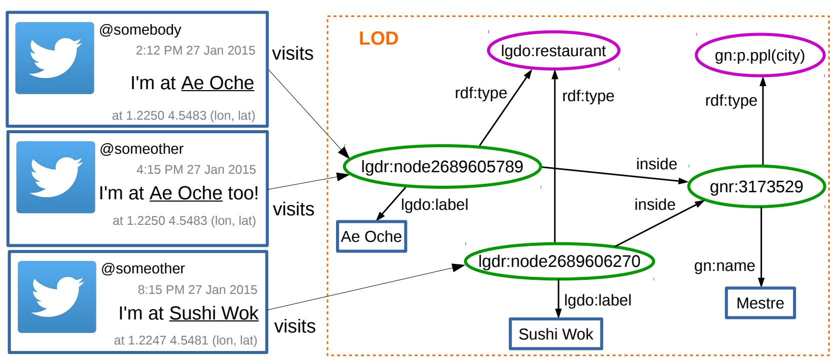

Fig. 1. Associations of tweets with PoIs and with their classes

Definition 3. *A* ***direct association*** *is a triple of the form* $(mds, property, r)$*, where* mds *is an MDS, property is the association kind, and* r *is a resource.*

Many associations can annotate the same MDS. Annotating movement segments that are sequences of positions instead of individual positions grants compactness (e.g., a move can be annotated as a whole with a transportation means).

Indirect associations of MDS with other resources can be derived from the direct associations by following properties between resources of LOD collections, as stated by Definition 4.

Definition 4. *An* ***indirect association*** *between a movement data segment* mds *and a resource* r *is a tuple* $(mds, path, r)$ *where path is a chain of links constituted by the direct association* $(mds, property, r')$ *of* mds *with the resource* r'*, followed by a sequence of LOD properties that semantically links* r' *to* r*.*

In other words, an indirect association $(mds, path, r)$ indirectly links the mds to a resource r via a $path$ composed by a direct association followed by some LOD property(ies). For example, the tweets on the left side of Fig. 1 are indirectly associated with the concept *restaurant* via direct *visits* associations followed by the *type* property between the respective resources and *restaurant*. They are also indirectly associated with the city of *Mestre* and its class (*gn:p.ppl (city)*).

Finally, a **Semantically-enriched Movement Dataset (SMoD)** is a set of movement data segments which are associated with resources for semantic annotation purposes. The use frequency of a resource in associations of an SMoD determines its number of hits with respect to that SMoD (Definition 5).

Definition 5. *Let $\mathcal{S}$ be a SMoD, $\mathcal{R}$ a set of resources, and $\mathcal{A}$ a set of direct and indirect associations between movement segments in $\mathcal{S}$ and resources in $\mathcal{R}$. The **number of hits** $h(r, \mathcal{A})$ of a resource $r \in \mathcal{R}$ with respect to A, is the number of distinct MDSs of $\mathcal{S}$ which are annotated with r via associations in $\mathcal{A}$.*

Of course, one could also count separately the number of direct and indirect hits of each resource, by considering only direct and indirect associations, respectively. In Figure 1, while the number of (direct) hits of the *lgdr:node2689605789* is 2 hits, its class *lgdo:Restaurant* has 3 (indirect) hits. Also the city of Mestre *gnr:3173529* and its class (*gn:p.ppl (city)*) have 3 (indirect) hits each one.

3 Approach to tailor Dimensions for MDWs

This section describes our proposal to automatically tailor dimensions from a given MoD. Figure 2 illustrates the steps of the Extract-Transform-Load (ETL) process that are required to build tailored dimensions, with those steps that are the focus of this work (2 and 3) in bold.

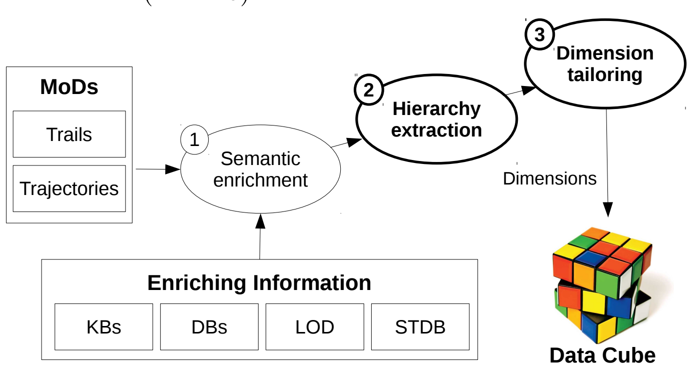

Fig. 2. The process for building semantic dimensions.

Step 1 (*Semantic enrichment*) takes as input a MoD and KBs, LOD, conventional databases (DBs), and/or spatio-temporal DBs (STDBs) containing information useful for making semantic annotations. As pointed out in [5], some important aspects to enrich MDSs are visited POIs, transportation means, environmental conditions, goals, and activities of the moving objects. This task is addressed in several works (e.g., [3, 11, 4]), but it is beyond the scope of this paper, whose focus in on steps 2 and 3.

Step 2 is called *Hierarchy extraction* and it starts from the MoD which has been semantically enriched. For each semantic annotation, we explore the resources which are directly associated with MDSs in order to find properties that connect these resources to other resources by allowing to extract a hierarchy of concepts or of objects. Sometimes the extracted hierarchies are not so detailed, e.g., some levels are missing, so to improve them we search for other representations of a resource in a different KB by using the property *owl:sameAs* and then we integrate several properties. For example, for spatial objects, we can merge properties from GeoNames and GADM to build a spatial hierarchy connecting POIs to the continent they belong to, and having as intermediate levels districts, cities, states and countries. Remember that resources can be hierarchically organized according to a variety of partial order relationships expressed by properties (e.g., *isA*, *partOf* or *contains*). It enables information analysis of data annotated with them at different levels of detail.

Step 3 is called *Dimension tailoring* and is aimed at creating views of the extracted hierarchies that can serve as tailored analysis dimensions for the given MoD. In fact, hierarchies obtained from step 2 may contain non-relevant elements for the actual data analysis, that can be filtered away. Hierarchies produced in this way can be quite smaller than the ones from which they are extracted, leading to efficiency gains, among other potential benefits.

The proposed algorithm (Algorithm 1) takes as input a hierarchy $\mathcal{H}$ from step 2, a set of associations $\mathcal{A}$ which annotate MDS with resources in $\mathcal{H}$ and a given threshold σ. Then it proceeds as follows:

Algorithm 1: Tailor($\mathcal{A}$, $\mathcal{H}$, σ)

Input: Set of associations $\mathcal{A}$, resource hierarchy $\mathcal{H}$, threshold σ
Output: dimension d

$d \leftarrow$ *new dimension*;
$Filter(\mathcal{H}.root, \sigma, \mathcal{A})$;
$Merge(\mathcal{H}.root)$;
$d.root \leftarrow FindNewRoot(\mathcal{H}.root)$;
return d;

Algorithm 2: Filter(r, σ, $\mathcal{A}$)

Input: hierarchy root r, threshold σ, set of associations $\mathcal{A}$
Output: none

if r *is a leaf* **then**
  $r.hits \leftarrow h(r, \mathcal{A})$;
else
  $r.hits \leftarrow 0$;
  for *each* $child \in r.children()$ **do**
    $Filter(child, \sigma, \mathcal{A})$;
    $r.hits \leftarrow r.hits + child.hits$;
if $r.hits < \sigma$ **then**
  $r.label \leftarrow$ "$Others$";

Algorithm 3: Merge(r)

Input: hierarchy root r
Output: none

$o \leftarrow$ *new node*;
$o.label \leftarrow$ "$Others$";
$o.hits \leftarrow 0$;
for *each* $child \in r.children()$ **do**
  if *child.label= "Others"* **then**
    $o.addChildren(child.children())$;
    $r.removeChild(child)$;
    $o.hits \leftarrow o.hits + child.hits$;
  else
    $Merge(child)$;
if $|o.children()| > 0$ **then**
  $r.addChild(o)$;
  $Merge(o)$;

Algorithm 4: FindNewRoot(r)

Input: hierarchy root r
Output: new hierarchy root

if $|r.children()| = 1$ **then**
  $child \leftarrow$ the unique child of r;
  return *FindNewRoot(child)*;
else
  return r

1. via a post-order traversal of $\mathcal{H}$, the function *Filter* directly computes the number of hits (Definition 5) in the leaves of $\mathcal{H}$ whereas for an internal node it calculates such a number by summing up the hits in its children. Moreover, it checks whether a node has a number of hits greater or equal than σ. If this constraint is not satisfied the label of this node is replaced by "Others", which is a dummy value that states that this node is not relevant for the dimension (Algorithm 2).
2. After the execution of *Filter*, an internal node could have several children labeled by the dummy value "Others", so the function *Merge* is aimed at replacing all these nodes with a single child having such a label (Algorithm 3).
3. Finally, the least common ancestor is computed and the hierarchy view rooted at this node is returned as dimension (Algorithm 4).

4 Experiments

In this section, we present the experiments conducted on two datasets of tweets which have been semantically enriched by using two criteria: spatial proximity of the MDS to a resource and textual similarity between the name of the resource and the content of the tweet [6].

The first dataset, **SMoD1**, contains 1,530 tweets which have been directly associated with 258 spatial resources by considering a spatial proximity of 8 meters and textual similarity of 94%. These tweets are indirectly associated with 129 spatial resource categories via 1,328 *rdf:type* assertions.

The second dataset, **SMoD2**, is composed by 10,710 tweets directly associated with 1,501 spatial resources by using less restrictive conditions: a spatial proximity of 16 meters and textual similarity of 90%. Tweets are indirectly associated with 338 spatial resource categories via 7,422 *rdf:type* assertions.

Both SMoD1 and SMoD2 have their tweets successfully associated with a unique spatial resource. Each tweet has a geo-location inside a Minimum Bounding Rectangle (MBR) within the Brazilian territory (NE 5.264860 -28.839041, SW -33.750702 -73.985527) and was extracted between June 9, 2014 and July 14, 2014. The Twitter extraction includes a condition to filter only tweets that start with the string "I'm at" and end with a FourSquare URL. Spatial resources inside the same Brazilian territory MBR were extracted from DBpedia and LGD on July 21, 2014 and they were considered for the semantic enrichment step.

In this paper, we assume that the spatial resources used as annotations in SMoD1 and SMoD2 represent the current place where a user sends the tweet at the annotated timestamp. Hence, we apply our method in order to build a spatial object dimension and a spatial category dimension for both SMoDs.

The spatial objects hierarchy was extracted by searching for DBPedia and LGD resources used as annotations in SMoD1 and SMoD2, and connecting them to the respective GeoNames and GADM resources via *sameAs* properties, as illustrated in Fig. 3. Then, for the found resources, the hierarchies are completed by exploiting properties like *gadm:in_country* and *gn:parentfeature* to connect them with upper level resources. A portion of the resulting hierarchy is shown in Fig. 4. It is worth noticing that the nodes are organized in various levels of abstraction: *idresource*, *district*, *city*, *state*, *country* and *continent*.

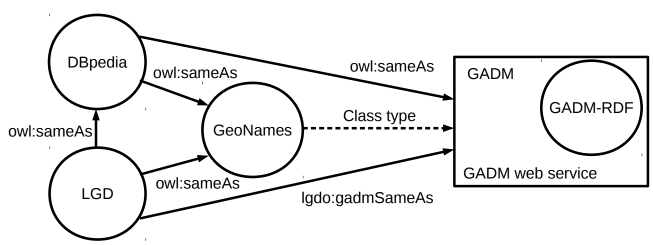

Fig. 3. LOD collections and properties used to extract hierarchies of spatial objects.

Algorithm 1 was used to create views of hierarchies such as the one illustrated in Fig. 4, in which the number of hits per node appears in brackets. Assuming a threshold $\sigma = 20$ groups of nodes at each level (e.g., *Santana*, *Mocca* and *Vila Mariana* in the level *district*) are replaced with the dummy value "Others'. This means that for the chosen threshold such resources are not so relevant since they annotate too few tweets. Hence they are filtered away. Figure 4 illustrates a spatial hierarchy, with a view composed of nodes having more than $\sigma = 20$

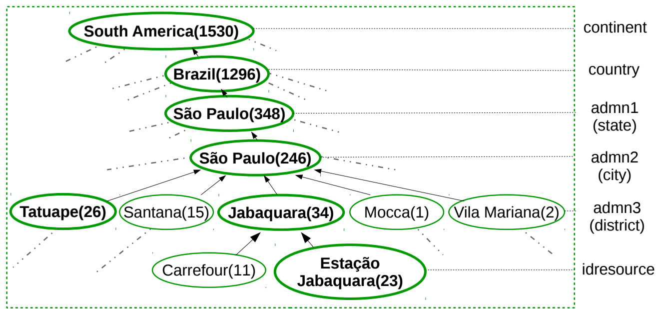

Fig. 4. A portion of a spatial objects hierarchy with the most popular nodes in bold.

hits in bold. Such a view added with a single dummy node condensing each set of siblings that have a lower number of hits (e.g., *Santana*, *Mocca* and *Vila Mariana*) can be used as a tailored dimension.

Figure 5 shows the number of resources (nodes) in each level of the tailored objects dimension generated for different values of the minimum number of hits per resource (σ) in SMoD1 (left) and SMoD2 (right). Notice that those numbers of dimension nodes falls sharply as σ grows. In addition, even for $\sigma = 1$ there is a considerable reduction in the number of countries (7 that originated the tweets in SMoD1, instead of around 200 in the world), cities (around 100 generating those tweets, instead of more than 6000 in the region defined by an MBR encompassing Brazil that was used to filter those tweets), just to mention some reductions.

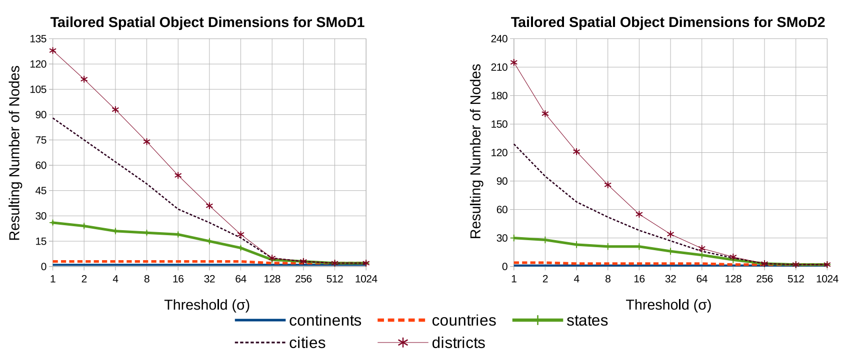

Fig. 5. Nodes per level of tailored spatial objects dimensions for increasing values of σ

A hierarchy of PoI concepts indirectly associated with the tweets in SMoD1 was built by exploiting *rdfs:subClassOf* properties between such concepts. Figure 6 shows an extract of a hierarchy of 93 concepts built from SMoD1. In

SMoD1, the most visited concepts are *Thing* (1509 hits), *Amenity* (981 hits), *Restaurant* (225 hits), *Station* (197 hits), *Shop* (156 hits), and *Cafe* (124 hits).

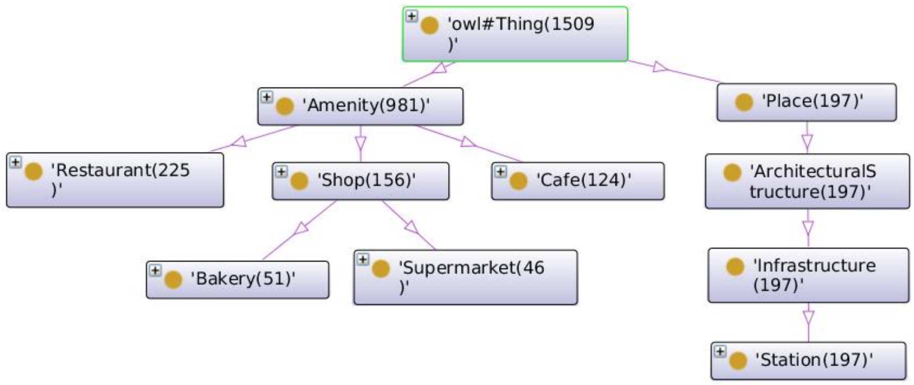

Fig. 6. LGD PoI categories with at least 15 indirect associations with tweets in SMoD1

5 Related Work

Several works propose models and approaches for analyzing movement data in DWs [12, 7, 8, 13, 14, 9]. Among them, only [9] accommodates MoD semantically annotated with LOD in a multidimensional model that separates hierarchies of objects and hierarchies of categories of these objects, and also allows hierarchies of movement segments, movement patterns, and their respective categories as dimensions. However, it does not provide means to build dimensions with LOD.

Many other works exploit semantic Web technology for building and operating on DWs [15–22]. Some of these proposals use semantic Web data such as LOD to feed DWs. [19] proposes a semi-automatic method for extracting semantic data into a multidimensional database on-demand. In [18] the authors follow a similar approach, but restrict their case study to statistical LOD and use the RDF Data Cube vocabulary[6]. [21] discusses modeling issues to support OLAP queries in SPARQL. Thus, these approaches foster the use of existent semantic Web data and tools.

Finally, some works also provide means for social media information analysis in DWs [23–25]. [24] proposes an interactive methodology for designing and maintaining Social Business Intelligence (SBI) applications. This methodology aims at providing quick responses to face the high dynamism of user generated contents. It is applied to case studies related to Italian politics and goods

[6] http://www.w3.org/TR/vocab-data-cube

consumption. [25] presents an approach based on meta-modeling for coupling dynamic and irregular topic hierarchies for analyzing social media data. However, the proposed architecture uses an ontology editor for organizing these hierarchies, instead of reusing available ones.

None of these works provide means to derive tailored dimensions for analyzing particular datasets, and more specifically movement datasets, from their semantic annotations. Our approach is the first to do so, by exploiting views of objects hierarchies and classes hierarchies, such as those that can be extracted from LOD collections currently available on the Web.

6 Conclusions and Future Work

This paper makes some advances in the construction of DWs for analyzing semantically enriched movement data (SMoD). It proposes an automatic approach for building tailored analysis dimensions from hierarchies of objects or classes used to semantically annotate the SMoD. The main contributions are: (i) a method to extract hierarchies of objects and classes from LOD by exploiting partial ordering relations present in available LOD, such as $part\_of$ and is_a; (ii) an algorithm to tailor dimensions for particular SMoD, based on views of the hierarchies of classes and objects used to semantically enrich the movement data; and (iii) a demonstration of the proposed method in a case study that generates spatial dimensions from hierarchies of PoIs and their classes used to semantically enrich social media user's trails composed of tweets.

The experience has taught us that it is usually easier to extract hierarchies of objects and classes from existing LOD collections than building them from scratch. Nevertheless, sometimes it is necessary to complement the semantic relations available in LOD collections by using other means (e.g., investigating containment relationships between geographic extensions of PoIs) to fulfill the absence of explicit relations between some pairs of resources. Experiments with our method to tailor dimensions in case studies with tweets semantically enriched with DBPedia and LGD resources showed considerable reductions in the size of the resulting dimensions compared to whole hierarchies extracted of LOD, even for low values of the hits count threshold (number of movement segments annotated per PoI or class of PoI).

This work, for the best of our knowledge, provides the first proposal for building DW dimensions from LOD used to semantically enrich movement data, and perhaps other kinds of data too. It constitutes an important step towards building DWs for semantics-enabled information analysis. Future work includes: (i) deeper investigation of the implications of variations of the proposed method for tailoring dimensions; (ii) experiments with other movement datasets enriched with different LOD collections and knowledge bases; (iii) inspection of the features of other LOD collections to improve the hierarchy extraction; (iv) further evaluation of the benefits of the proposed method.

7 Acknowledgments

This work was supported by the European Union IRSES-SEEK (grant 295179), CNPq (grant 478634/2011-0), CAPES, and FEESC, the MIUR Project PON ADAPT (no. SCN00447), and by MOTUS (no. MS01 00015 - Industria2015). Special thanks to Cleto May and Douglas Klein for providing semantically enriched data for our experiments.

References

1. Parent, C., Spaccapietra, S., Renso, C., Andrienko, G.L., Andrienko, N.V., Bogorny, V., Damiani, M.L., Gkoulalas-Divanis, A., de Macêdo, J.A.F., Pelekis, N., Theodoridis, Y., Yan, Z.: Semantic trajectories modeling and analysis. ACM Comput. Surv. **45** (2013) Article 42.
2. Pelekis, N., Theodoridis, Y.: Mobility Data Management and Exploration. Springer (2014)
3. Yan, Z., Chakraborty, D., Parent, C., Spaccapietra, S., Aberer, K.: Semantic trajectories: Mobility data computation and annotation. ACM TIST **4** (2013)
4. Fileto, R., Krüger, M., Pelekis, N., Theodoridis, Y., Renso, C.: Baquara: A Holistic Ontological Framework for Movement Analysis Using Linked Data. In: ER. Volume 8217 of LNCS., Springer (2013) 342–355
5. Bogorny, V., Renso, C., de Aquino, A.R., de Lucca Siqueira, F., Alvares, L.O.: CONSTAnT - A Conceptual Data Model for Semantic Trajectories of Moving Objects. T. GIS **18** (2014) 66–88
6. May, C., Fileto, R.: Connecting Textually Annotated Movement Data with Linked Data. In: IX Regional School on Databases. ERBD, São Francisco do Sul, SC, Brazil (in Portuguese), SBC (2014)
7. Raffaetà, A., Leonardi, L., Marketos, G., Andrienko, G., Andrienko, N., Frentzos, E., Giatrakos, N., Orlando, S., Pelekis, N., Roncato, A., Silvestri, C.: Visual Mobility Analysis using T-Warehouse. IJDWM **7** (2011) 1–23
8. Wagner, R., de Macêdo, J.A.F., Raffaetà, A., Renso, C., Roncato, A., Trasarti, R.: Mob-Warehouse: A semantic approach for mobility analysis with a Trajectory Data Warehouse. In: SecoGIS 2013. Volume 8697 of LNCS. (2014) 127–136
9. Fileto, R., Raffaetà, A., Roncato, A., Sacenti, J.A.P., May, C., Klein, D.: A semantic model for movement data warehouses. In: DOLAP 2014. (2014) 47–56
10. Ngomo, A.N., Auer, S., Lehmann, J., Zaveri, A.: Introduction to linked data and its lifecycle on the web. In: 10th Intl. Summer School on Reasoning on the Web in the Big Data Era, Athens, Greece. Volume 8714 of LNCS. (2014) 1–99
11. Rinzivillo, S., de Lucca Siqueira, F., Gabrielli, L., Renso, C., Bogorny, V.: Where Have You Been Today? Annotating Trajectories with DayTag. In: SSTD. Volume 8098 of LNCS., Springer (2013) 467–471
12. Leal, B., de Macêdo, J.A.F., Times, V.C., Casanova, M.A., Vidal, V.M.P., de Carvalho, M.T.M.: From Conceptual Modeling to Logical Representation of Trajectories in DBMS-OR and DW Systems. JIDM **2** (2011) 463–478
13. Pelekis, N., Theodoridis, Y., Janssens, D.: On the management and analysis of our lifesteps. SIGKDD Explorations **15** (2013) 23–32
14. Leonardi, L., Orlando, S., Raffaetà, A., Roncato, A., Silvestri, C., Andrienko, G.L., Andrienko, N.V.: A general framework for trajectory data warehousing and visual OLAP. GeoInformatica **18** (2014) 273–312

15. Sell, D., Cabral, L., Motta, E., Domingue, J., dos Santos Pacheco, R.C.: Adding semantics to business intelligence. In: DEXA, Copenhagen, Denmark, IEEE Computer Society (2005) 543–547
16. Filho, S.I.V., Fileto, R., Furtado, A.S., Guembarovski, R.H.: Towards Intelligent Analysis of Complex Networks in Spatial Data Warehouses. [26] 134–145
17. Deggau, R., Fileto, R., Pereira, D., Merino, E.: Interacting with spatial data warehouses through semantic descriptions. [26] 122–133
18. Kämpgen, B., Harth, A.: Transforming statistical linked data for use in OLAP systems. In: Proceedings of the 7th International Conference on Semantic Systems. I-Semantics '11, New York, NY, USA, ACM (2011) 33–40
19. Nebot, V., Llavori, R.B.: Building data warehouses with semantic web data. Decision Support Systems **52** (2012) 853–868
20. Ibragimov, D., Hose, K., Pedersen, T.B., Zimányi, E.: Towards exploratory OLAP over linked open data: A case study. In: 7th Intl. Workshop on Business Intelligence for the Real-Time Enterprises (BIRTE 2014), Springer, Springer (2015) To appear.
21. Etcheverry, L., Vaisman, A.A., Zimányi, E.: Modeling and querying data warehouses on the semantic web using QB4OLAP. In: DaWaK. Volume 8646 of LNCS. (2014) 45–56
22. Abelló, A., Romero, O., Pedersen, T.B., Llavori, R.B., Nebot, V., Cabo, M.J.A., Simitsis, A.: Using semantic web technologies for exploratory OLAP: A survey. IEEE Trans. Knowl. Data Eng. **27** (2015) 571–588
23. Llavori, R.B., Aramburu, M.J., Llidó, D.M., García-Moya, L.: Towards a semantic data infrastructure for social business intelligence. In: ADBIS. Volume 241 of Advances in Intelligent Systems and Computing., Springer (2013) 319–327
24. Francia, M., Golfarelli, M., Rizzi, S.: A methodology for social BI. In: 18th Intl. Database Engineering & Applications Symposium, (IDEAS), Porto, Portugal, ACM (2014) 207–216
25. Gallinucci, E., Golfarelli, M., Rizzi, S.: Meta-stars: Dynamic, schemaless, and semantically-rich topic hierarchies in social BI. In: 18th Intl. Conf. on Extending Database Technology, EDBT 2015, Brussels, OpenProceedings.org (2015) 529–532
26. Bogorny, V., Vinhas, L., eds.: XI Brazilian Symposium on Geoinformatics, November 28 to December 01, 2010, Campos do Jordão, São Paulo, Brazil. In Bogorny, V., Vinhas, L., eds.: GeoInfo, MCT/INPE (2010)