

Cross-domain Semantic Drift Measurement in Ontologies Using the SemaDrift Tool and Metrics

Thanos G. Stavropoulos¹, Efstratios Kontopoulos¹, Albert Meroño Peñuela², Stavros Tachos¹, Stelios Andreadis¹, Ioannis Kompatsiaris¹

¹Information Technologies Institute, Thessaloniki, Greece
{athstavr, skontopo, staxos, andreadisst, ikom}@iti.gr

²Vrije Universiteit Amsterdam, The Netherlands
albert.merono@vu.nl

Abstract. Detecting and measuring semantic drift in different versions of ontologies across time is a novel area of research that rapidly gains attention. Nevertheless, there exist only a few relevant practical methods and tools and even fewer are flexible enough to be efficiently applied to multiple domains. As the often domain-specific nature of ontologies may render methods and tools for measuring semantic drift ineffective, this paper presents the application and findings of the SemaDrift suite of methods and tools in several domains, illustrating novel insights for the first time. While developed in the context of the PERICLES FP7 project, aimed at Digital Preservation, domain-independent text and structural similarity measures, available both as a software library and as a Protégé plugin for end-users, are now applied in the Dutch Historical Census and the BBC Sports Ontology. The two different domains demonstrate its applicability and ability to pinpoint the location, nature, origins and destinations of concept drift.

Keywords: Semantic drift; concept drift; semantic change; ontologies; Protégé.

1 Introduction

As the world continuously changes, concepts and their underlying semantics also change over time. In digital environments that are consequently also subject to continual change, ensuring that digital content remains understandable – and from this perspective accessible and reusable – poses a formidable challenge. The evolution of semantics is an active area of research, especially challenged by the lack of universal metrics to address specificities and peculiarities pertinent to each domain.

Evolving semantics, also referred to as *semantic change*, observes and measures the phenomenon of change in the meaning of concepts within knowledge representation models, along with their potential replacement by other meanings over time. In the Semantic Web [1], the representation of the underlying knowledge is typically assumed by ontologies. Thus, it can be easily perceived that semantic change can have drastic consequences on the use of ontologies in Semantic Web and Linked Data applications. In this setting, semantic change, i.e. the structural difference of the same

concept in two ontologies [2], relates to various lines of research. Such examples are *concept* and *topic shift* [3], *concept change* [4], *semantic decay* [5], *ontology versioning* [6] and *evolution* [7]. A brief disambiguation of these terms can be found in [8].

Semantic drift can be defined as the phenomenon of ontology concepts gradually changing as our knowledge of the world evolves, obtaining possibly different meanings, as interpreted by various user communities or in a different context, risking their rhetorical, descriptive and applicative power [9]. *Concept drift* can refer to this language-related phenomenon, but also in abrupt parameter value changes in data mining [10].

This paper presents findings in two vastly diverse domains through applying a novel set of universal, domain agnostic semantic drift metrics across various domains using the SemaDrift suite of tools and metrics. The metrics, initially presented in [8], are embedded in respective software tools, that offer the means for domain experts to assess drift without programming knowledge. Namely, the SemaDrift plugin for the Protégé platform¹ aims at assisting a wider audience to monitor and manage concept drift and was developed in the context of the PERICLES FP7 project², integrating and extending existing studies [3] and previously developed open, reusable methods [8], [11].

The domains studied in this paper are (a) *the CEDAR dataset* containing historical Dutch census data, and (b) *the BBC Sport Ontology* for representing competitive sports events. In the historical census domain, the metrics help pinpoint historical occupation qualities for the population between 1869 and 1930, almost on the fly. Most importantly, using the same tool we move on to the BBC Sport Ontology where the metrics pinpoint the location, nature, origins and destinations of concept drift, across six versions. The ontologies used in this work and the SemaDrift outputs are publicly available at <https://github.com/skontopo/MEPDaW2017>.

The rest of the paper is structured as follows: Section 2 presents related work in metrics and tools for measuring drift. Section 3 presents the underlying metrics and the SemaDrift framework. Sections 4 and 5 present the two proof-of-concept scenarios and report on our findings. Conclusions and directions for future work are listed in the final section.

2 Related Work

Measures of semantic richness of Linked Data concepts have been investigated in [5], proving that increasing reuse of concepts decreases its semantic richness. Other studies have examined change detection between two ontologies at a structural or content level [2], while ontology mapping investigates the relationships and correspondences between two ontologies [12]. Concept drift has been measured either by clustering while populating ontologies [13] or by applying linguistic techniques on textual concept descriptions [14]. A vector space model by random indexing has been utilized to track changes in an evolving text collection [9] and to visualize the drift of vocabular-

¹ The Protégé Ontology Editor: <http://protege.stanford.edu>

² PERICLES FP7 project: www.pericles-project.eu

ies in a diachronic sample of the Linked Open Data cloud [15]. A strategy to represent change has been based on ontology evolution [7]. However, most of these techniques are not directly applicable to Semantic Web constructs or present limited statistical data.

An appealing solution we have adopted transfers the notions of *label*, *extension* and *intension* from machine learning concept drift to semantic drift, further defining them in ontology terms [3]. A key reason for choosing the specific approach lies in the fact that the authors of [3] consider both linguistic aspects as well as the structure of the ontology itself, which to the best of our knowledge is the most complete methodology we came across.

Much philosophical debate examines how and by which properties a concept can be identified across time and appropriate formalization [16]. Some have utilized the notions of *perdurant* and *endurance* [17], so as to seek identity, by defining rigid properties that have to be persistent across instances and, thus, can identify entities [10]. In this work, we adopt, implement and integrate the methods in [3] into a familiar application for knowledge engineers, targeting not only the lack of reproducible cross-domain metrics for semantic drift but also the lack of similar graphical user interfaces.

3 Semantic Drift Metrics and the SemaDrift Platform

3.1 Semantic Drift Metrics

The drift metrics considered here implement and extend previous work in the field of concept drift ([8] and [11]), where highly applicable notions and metrics for measuring concept drift in the context of data mining have successfully been transferred to semantic drift. The method to measure concept drift in semantics considers two basic factors: (a) the different aspects of change and (b) whether concept identity is known or not. The aspects of change can be:

- *Label*, which refers to the description of a concept, via its name or title;
- *Intension*, which refers to the characteristics implied by it, via its properties;
- *Extension*, which refers to the set of things it extends to, via its number of instances.

Meanwhile, the correspondence of a concept across versions can be either known or unknown, resulting in two different approaches for measuring change:

- *Identity-based approach* (i.e. known concept identity): Assessing the extent of shift or stability of a concept's meaning is performed under the assumption that its identity is known across ontologies. For instance, considering an ontology *A*, and its evolution, ontology *B*, each concept of *A* is known to correspond to a single, known concept of *B*.
- *Morphing-based approach* (i.e. unknown concept identity): Each concept is pertaining to just a single moment in time (ontology), while its identity is unknown

across versions (ontologies), as it constantly evolves/morphs into new, even highly similar, concepts. Therefore, its change has to be measured in comparison to every concept of an evolved ontology.

The currently proposed method considers the more general morphing-based approach and considers drift as the dissimilarity of two maximally similar concepts in two versions [3]. Despite several methods have been proposed to seek identity correspondence across versions [10], they still can be domain- or model-dependent, mandating for ad-hoc expert knowledge in the form of annotations, user input or using explicit identities. In order to measure change, the meaning of each concept at a given point in time is defined as a set of the three different aspects, as follows:

$$C^t = \langle label_t(C), int_t(C), ext_t(C) \rangle$$

where C^t denotes the meaning of concept C at point t . Each of its aspects, $label_t(C)$, $int_t(C)$ for intensional and $ext_t(C)$ for extensional, is measured as follows:

$$label_t(C) = \{l, \langle C, rdfs: label, l \rangle \in T\}$$

$$int_t(C) = \{i, i = \langle P, x, C \rangle, x = rdfs: domain \vee x = rdfs: range, \forall i \in T\}$$

$$ext_t(C) = \{x, \langle x, rdf: type, C \rangle \in T\}$$

where T is the set of all triples in version t of the ontology. More concretely:

- *label* is the *rdfs: label* of a concept (a string);
- *intension* is a set of triples (i.e. the properties that involve the concept, calculated as the union of all RDF triples with C in the subject or object position of OWL object or datatype properties;
- *extension* is the set of strings (i.e. the names of instances with the concept as value of *rdf: type*).

Due to the morphing based approach, each concept's drift is measured as the average drift to all concepts of the next ontology. Comparisons for strings are made using the Monge-Elkan algorithm [18], found to optimally suit strings in ontologies such as CamelCase or snake_case, and Jaccard similarity for sets.

In detail, if n_2 is the total number of concepts in t_2 , we define label, intensional and extensional drifts of C between versions t_1 and t_2 as follows:

$$label_{t_1 \rightarrow t_2}(C) = \frac{\sum_{i=1}^{n_2} MongeElkan(label_{t_1}(C), label_{t_2}(C_i))}{n_2}$$

$$int_{t_1 \rightarrow t_2}(C) = \frac{\sum_{i=1}^{n_2} Jaccard(int_{t_1}(C), int_{t_2}(C_i))}{n_2}$$

$$ext_{t_1 \rightarrow t_2}(C) = \frac{\sum_{i=1}^{n_2} Jaccard(ext_{t_1}(C), ext_{t_2}(C_i))}{n_2}$$

As all metrics have a range of [0, 1], their average (without using weights in the general case) can be considered as an overall, *whole* aspect:

$$whole_{t_1 \rightarrow t_2}(C) = \frac{label_{t_1 \rightarrow t_2}(C) + int_{t_1 \rightarrow t_2}(C) + ext_{t_1 \rightarrow t_2}(C)}{3}$$

3.2 The SemaDrift Tool to Measure Semantic Drift

While the SemaDrift metrics discussed above can be used directly in third-party scripts and software via the SemaDrift API Library, a domain expert or researcher may not possess the ability to do so. Especially in the case that several domain ontologies need to be explored, as in this study, researchers need to use a common tool to work fast, reliably and without further adaptations as a common point of reference. The SemaDrift Protégé plugin serves this purpose³.

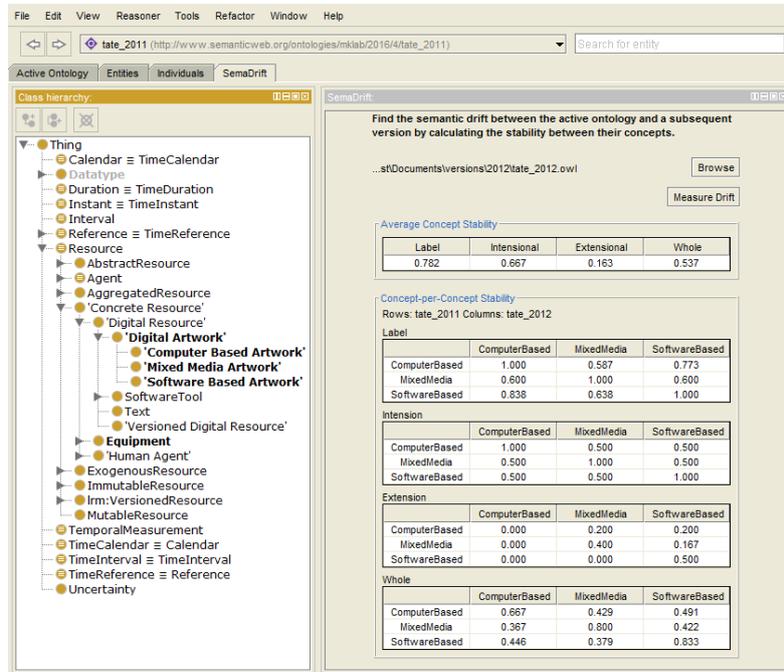


Fig. 1. SemaDrift Protégé plugin: The native tree hierarchy of the open ontology is shown on the left, while the plugin-provided content resides on the right, showing a second ontology to compare to, accompanied by the respective measurements.

The plugin's main panel is shown in Fig. 1. The tool provides a subset of the basic functions of the underlying SemaDrift API in a graphical manner. For that purpose, it

³ The SemaDrift suite is available at: <http://mklab.itι.gr/project/semadrift-measure-semantic-drift-ontologies>.

exposes some of its functions and accommodates the outcomes in suitable user controls using the Java Swing library. This edition of the plugin focuses on ontology pairs, i.e. two versions of the same ontology, in order to provide more insight into them and their differences, fitting also into the Protégé workspace philosophy. Usually, the users work on a single ontology at a time, which is always displayed as a tree hierarchy of classes at the left pane. Then, plugins occupy the right pane, which is free to accommodate their functions (Fig. 1).

As a first step, the user has to select the pair of ontologies for which to measure drift. To take advantage of the environment, the plugin assumes that the first selected ontology is the one currently loaded in Protégé, allowing also its in-depth visualization, reasoning and query execution. The second ontology can be selected from the SemaDrift pane using the “Browse” button to look through local or remote storage.

After both ontologies are available, pressing on the “*Measure Drift*” button will display the SemaDrift metric results. Stability, as a measure of drift, is shown in two sections: overall average stability per aspect and concept pair stability for all aspects. The first section constitutes the most generic, abstract measure of drift. It displays a table with the average drift of all concepts from the former ontology to the latter, per each of the four aspects: label, intension, extension, and whole. Naturally, the measurements are derived using the metrics and algorithms for each aspect described in the previous section, yielding a value from zero (no similarity) to one (full similarity).

The second section of results is displayed in respective tables. Each table row corresponds to a concept of the former ontology and each column to a concept of the latter. Consequently, each cell holds the similarity metric (i.e. concept stability) between each pair of concepts. These similarity values between pairs can further be utilized by users for different purposes for example to generate similarity graphs or *morphing chains* such as those demonstrated in the rest of the paper.

4 Dutch Historical Censuses

Census data are essentially time series of systematic population records, and hence an important source for studying semantic drift of concepts involving culture and economics. In the Netherlands, the Dutch historical censuses are 17 country-wide population reports performed between 1795 and 1971, once every 10 years. In each of these reports, the government counted the population of the country and its demographic, occupational, and housing characteristics. In 1971, this detailed reporting stopped due to social concern on privacy. Nevertheless, the exhaustive, detailed, and aggregated characteristics⁴ of these censuses have continued to attract the attention of historians and social scientists [19], who nowadays study them via a collection of 507 machine-readable spreadsheets, containing 2,288 census tables⁵. In order to improve more systematic and universal access to reproduce results of studies on this dataset, recent

⁴ Notably, the microdata registers (i.e. individual survey data), upon which these censuses are built, have been lost over time. Hence, the numbers are only aggregations, with no tracking information leading to the original individuals.

⁵ See <http://volkstellingen.nl/>

efforts have managed to publish it as Linked Data (LD) [20]. This LD set of the Dutch historical census will be further called the CEDAR (Center of Excellence for Document Analysis and Recognition) dataset, after its creators.

SemaDrift was used to analyze semantic drift in the CEDAR dataset and gain insights as to how *occupational concepts*, describing citizen jobs, changed in the period 1869-1930. For each of these two years, these occupational concepts are described with three attributes: (a) the *occupational concepts* themselves are SKOS concepts [21] using URIs of the Historical International Standard Classification of Occupations [22] (HISCO); (b) the *number of persons* having a specific job are associated with those SKOS concepts; and (c) a *number of labels* (in Dutch) are associated with these SKOS concepts. This implies that the *constraints* inherent to these data with respect to *intensions*, *extensions*, and *labels* are as follows: intensions do not exist for these occupational concepts, since no further formal descriptions are available; extensions are restricted by the cardinality of the concepts; and labels are assigned and abundant for all concepts in both years. An example is shown in Fig. 2.

A data transformation step is required before feeding the dataset into the tool to address not only the format but also to generate more meaningful properties. Namely, the data format, as originally shown in Fig. 2, is transformed to two OWL ontologies for the years 1869 and 1930. To do so, we convert every HISCO `skos:Concept` to an `owl:Class`, assigning them all `rdfs:label` in the original data. Furthermore, to obtain an accurate representation of extensions, we unroll the integer counts, as seen in Fig. 2, and generate as many anonymous instances as specified by these integers. This is done since, following a proper ontological representation, the extensional aspect actually refers to instances and not numerical properties. Finally, we assign these anonymous instances to their corresponding HISCO `owl:Class` using an `rdf:type` relation, thus each of them representing one person that carried the job indicated by the class.

```
cedar:hisco-13100 a skos:Concept ;
    skos:prefLabel "Hoogleeraar"@nl .

cedar:BRT-1930/37860 a qb:Observation ;
    cedar:occupation cedar:hisco-13100 ;
    qb:dataSet cedar:BRT-1930 ;
    cedar:populationSize "52"^^xsd:integer .
```

Fig. 2. Excerpt of the original CEDAR data. Census counts are modeled as RDF Data Cube observations, which carry information about the occupation class and the number of persons belonging to it.

After using SemaDrift for the two ontologies, respective average per aspect stability and average concept-per-concept stability are generated. The latter is used to draw morphing chains for topics of interest. After observing the table, first, the most stable concept between both versions is the occupational class `hisco:-1` (stability of 0.917 – see Fig. 3), the class for occupations that cannot be classified elsewhere in HISCO. This is due to a great label (0.750) and extensional (1.000) stability, which suggests

that both the coherency of data coders w.r.t. unclassifiable jobs, and the population carrying those remained stable in this period.

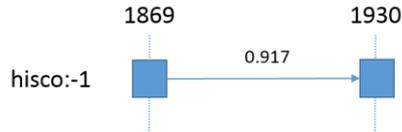


Fig. 3. Morphing chain of the whole aspect of the `hisco:-1` concept.

According to previous studies in extensional semantic drift in this dataset [23], other interesting classes from 1869 with expected extensional drift are:

- `hisco:97125`, loaders of ships, trucks, wagons or airplanes. These workers do not appear again in 1930, and the stability w.r.t. similar classes, like `hisco:97145` (storehouse workers), is significantly lower (0.479). Their closest matches in terms of stability are varnishers and stone polishers (0.717);
- `hisco:21110`, general managers. This group does appear in 1930, but the similarity of their classes has greatly drifted (0.511). Many other occupational jobs, with loose semantic similarity, display more stability w.r.t. the original class;
- `hisco:41025`, working proprietors. Similarly, this group of workers shows a great deal of drift, to the extent of not having an equivalent class in 1930. This might be due to historical reasons, i.e. the late industrialization in the Netherlands and its effects on evolving old small business owners into upper-class company investors. Noticeably, the related class `hisco:43200`, commercial agents, displays certain stability (0.405).

According to these results, we can extract various conclusions from the evolution of occupational concepts in the CEDAR dataset. We notice that the metrics shown in an extension drift analysis strictly based on concept cardinalities are mostly useful when an underlying stable schema that holds over time is previously given (in our case, HISCO). These metrics can be used to better assess the quality of expert-curated annotations and mappings (e.g. from strings to HISCO codes), but also to evaluate to what extent these mappings can be reused in different conditions. In this respect, the SemaDrift results can be used to evaluate the time period for which a certain ontology class can be safely used to access database content that evolves over time.

It is important to underline some intrinsic limitations in the study of semantic drift within the CEDAR dataset. Besides the lack of concept intensions, many of these limitations are related to the problem of *identity*, as also reported by [3]. First, the identity between *classes* cannot be assumed even between those of identical HISCO codes, since these are convoluted and culturally changing time periods. Secondly, the identity between *instances* of these classes is even more volatile. Human annotated identity information such as the existence of `owl:sameAs` links between instances of different time periods would greatly improve the outcomes of the extensional drift analysis, but require manual labor. Finally, using class cardinalities as a proxy for $ext_t(C)$ is a natural limitation of the CEDAR dataset. However, we use these cardi-

nalities: (a) to show that $ext_t(C)$ can be flexibly adapted to scenarios where instance data is limited; (b) coherently with classic measures of ontology evolution [7], e.g. leveraging the fact that concepts with a decreasing number of instances over time are likely to disappear or be merged with others; and concepts with an increasing number of instances are likely to be split and/or specialized.

The initial data transformation effort required in this scenario is justified, as the tool assumes proper ontology format (OWL) and design (instances instead of numeric properties according to their meaning). However, the tool itself with the existing morphing-based metrics is apparently very useful to very quickly gain access to insights regarding the evolution of semantic concepts that would otherwise require serious labor.

5 The BBC Sport Ontology

BBC is one of the pioneers in the field of ontology-based technologies, using them at an industrial level since 2010⁶. In the past, they found that conventional content management systems impose serious limitations on the flexibility of the ways that content is served, limiting the richness of the experience they offer to their visitors. To overcome these limitations and enhance the experience for website users, they turned to ontologies and Linked Data. An additional key benefit is that this approach also significantly reduces the time it takes for editors to create content that is easily discoverable across the website.

One of the first ontologies developed by BBC was the *Sport Ontology*⁷, which initially started as an effort to represent information about the competitions, teams, players and matches of the 2010 World Cup. However, although it originated as a specific use case, the Sport Ontology has since been significantly extended and is now applicable to representing a wide range of competitive sporting events. The BBC now use this ontology to support their sports coverage, including coverage of both the 2012 London Olympics and the 2014 Brazil World Cup.

Table 1. Information regarding the BBC Sport Ontology versions studied in the paper.

Ontology version	v2.10	v2.11	v2.12	v2.13	v3.0	v3.2
Date created	20-Feb-14	27-Mar-14	16-Sep-14	9-Feb-15	8-Apr-15	14-May-15
Class count	35	37	37	37	39	38
Property count	50	49	49	49	49	49
Individual count	27	21	33	41	41	41

⁶ BBC ontologies homepage: <http://www.bbc.co.uk/ontologies>

⁷ BBC Sport Ontology homepage: <http://www.bbc.co.uk/ontologies/sport>

The ontology’s significance for BBC, along with its potential applicability in various sports-related deployments, has led to our inclusion of the Sport Ontology in this study. However, compared to the case study presented in the previous subsection, the scope of this analysis is different, in the sense that we are investigating design decisions from version to version, possibly influenced by the company’s intended end-user applications and the public’s demonstrated preference to certain pertinent aspects. Table 1 contains information regarding the versions of the Sport Ontology studied in this paper⁸.

The six different versions of the Sport Ontology were loaded in SemaDrift. As derived by the tool (but also implied in the table), the ontology is extremely stable with regards to its intensional aspect (i.e. classes and properties), with most classes demonstrating a perfect stability of 1. This implies that the set of concepts included in the ontology are almost finalized, and the ontology itself has matured enough, rarely undergoing significant modifications in its structure. An exception was observed for classes `CompetitiveSportingGroup` and `CompetitiveSportingOrganisation`, whose stabilities were reported at an average of 0.9 each, due to changes in domains and ranges of respective properties in versions 2.11 and 2.12. After consulting the ontology documentation, we deduced that these changes coincide with some corrections to the corresponding properties introduced by BBC’s ontology engineers.

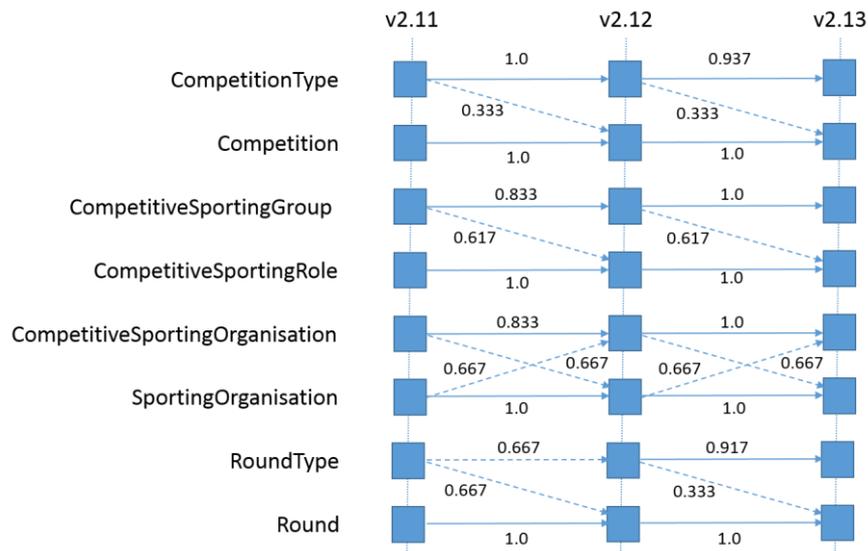


Fig. 4. Morphing chains illustrating the drifts of the “whole” aspect in versions 2.11, 2.12 and 2.13 of the BBC Sport Ontology.

On the other hand, the ontology is less stable extensionally, which is mostly due to instances being added to specific classes in versions 2.11, 2.12 and 2.13, indicating

⁸ Note that versions prior to v2.10 were not available online on the BBC website.

that the specific versions of the ontology underwent population by certain individuals. More specifically, the initially empty (i.e. no instances) class `RoundType` was populated with 12 instances in version 2.12 (e.g. `final`, `quarter-final`, `semi-final` etc.) and with 4 additional instances in version 2.13. Additionally, class `CompetitionType`, which initially had 17 instances (e.g. `domestic-cup`, `european-cup`, `international` etc.), was populated with 4 additional instances in version 2.13. Overall, version 2.13 was the one where the extension of the ontology was finalized. Fig. 4 illustrates the relevant morphing chains illustrating the drifts of the “whole” aspect for the three ontology versions involving the respective classes. Finally, the changes in the two most recent versions of the ontology (3.0 and 3.2) were minimal, thus indicating that the ontology has been eventually stabilized.

Conclusively, the above study offers an insight into the design choices of the six most recent versions of the BBC Sport Ontology, e.g. corrections in the schema, population with additional types of sporting competitions etc. Admittedly, the investigations would probably be more intriguing if we could study earlier versions of the ontology, but unfortunately we were unable to retrieve the latter at the time of preparing this work.

6 Conclusions and Future Work

This paper employed novel ways to measure semantic drift in two different domains: historical censuses and competitive sports events. SemaDrift is a suite of tools, a software library, and an application, that can measure drift aspects for ontologies on-the-fly. Linked data for the Dutch historical census from 1869 to 1930 were transformed to OWL and processed to show interesting insights for the semantic change in the population’s occupation concepts. Moreover, semantic drift was studied for six versions of the BBC Sport Ontology. Using the same tool to gain insights for unrelated domains demonstrated its universal and cross-domain properties. Also, its usefulness is shown, as it gives access to insights otherwise hard to obtain, such as to assess the nature of the drift (extensional), locate it in time and track the migration of meaning from concept to concept through morphing chains.

Future work will be focused on expanding to more domains and extending the tools. As already apparent in this study, the tool may handle most ontologies out-of-the-box enabling researchers without programming knowledge to do more. However, the historical censuses have uncovered not only a minor change in format but also in ontology design. As both these matters were solved by writing a transformation script, such scripts may be incorporated into the tool for future use. Furthermore, the lack of matching identities, elaborated on in previous studies [8], may be handled by alternative metrics. Future efforts could help define a gold standard to compare approaches in terms of precision and recall. Finally, additions to the tool’s GUI include handling more ontologies, adding visual aids and drawing abilities for morphing chains evolving the tool into a one-stop-shop for semantic drift measurement.

Acknowledgments. This research received funding from the European Commission Seventh Framework Programme under Grant Agreement Number FP7-601138 PERICLES.

References

1. Berners-Lee, T., Hendler, J., Lassila, O.: The semantic web. *Sci. Am.* 284, 28–37 (2001).
2. Tury, M., Bielíková, M.: An approach to detection ontology changes. In: Workshop proceedings of the sixth international conference on Web engineering - ICWE '06. p. 14. ACM (2006).
3. Wang, S., Schlobach, S., Klein, M.: Concept drift and how to identify it. *J. Web Semant.* 9, 247–265 (2011).
4. Uschold, M.: Creating, integrating and maintaining local and global ontologies. In: Proceedings of the First Workshop on Ontology Learning (OL-2000) in conjunction with the 14th European Conference on Artificial Intelligence (ECAI 2000), Berling, Germany.
5. Pareti, P., Klein, E., Barker, A.: A Linked Data Scalability Challenge: Concept Reuse Leads to Semantic Decay. In: Proceedings of the ACM Web Science Conference. ACM Press-Association for Computing Machinery (2016).
6. Yildiz, B.: Ontology Evolution and Versioning. Tech. Report, TU Vienna. (2006).
7. Stojanovic, L., Maedche, A., Motik, B., Stojanovic, N.: User-driven ontology evolution management. *Knowl. Eng. Knowl. Manag. Ontol. Semant. Web.* 133–140 (2002).
8. Stavropoulos, T.G., Andreadis, S., Riga, M., Kontopoulos, E., Mitzias, P., Kompatsiaris, I.: A Framework for Measuring Semantic Drift in Ontologies. In: 1st Int. Workshop on Semantic Change & Evolving Semantics (SuCESS'16). CEUR Workshop Proceedings, Leipzig, Germany (2016).
9. Wittek, P., Darányi, S., Kontopoulos, E., Moysiadis, T., Kompatsiaris, I.: Monitoring term drift based on semantic consistency in an evolving vector field. In: 2015 International Joint Conference on Neural Networks (IJCNN). pp. 1–8. IEEE, Killarney, Ireland (2015).
10. Meroño-Peñuela, A., Hoekstra, R.: What is Linked Historical Data? In: Proceedings of the 19th International Conference on Knowledge Engineering and Knowledge Management (EKAW 2014). pp. 282–287. Springer International Publishing (2014).
11. Stavropoulos, T.G., Andreadis, S., Kontopoulos, E., Riga, M., Mitzias, P., Kompatsiaris, I.: SemaDrift: A Protégé Plugin for Measuring Semantic Drift in Ontologies. In: Hollink, L., Darányi, S., Meroño Peñuela, A., and Kontopoulos, E. (eds.) 1st International Workshop on Detection, Representation and Management of Concept Drift in Linked Open Data (Drift-a-LOD) in conjunction with the 20th International Conference on Knowledge Engineering and Knowledge Management (EKAW). pp. 34–41. CEUR Workshop Proceedings Vol 1799, Bologna, Italy (2016).
12. Kalfoglou, Y., Schorlemmer, M.: Ontology mapping: the state of the art. *Knowl. Eng.* (2003).
13. Fanizzi, N., Amato, C., Esposito, F.: Conceptual Clustering: Concept Formation, Drift and Novelty Detection. In: The Semantic Web: Research and Applications, 5th European Semantic Web Conference, ESWC 2008, Tenerife, Canary Islands, Spain, June 1-5, 2008, Proceedings. pp. 318–332. Springer (2008).

14. Gulla, J., Solskinnsbakk, G., Myrseth, P.: Semantic Drift in Ontologies. In: Proceedings of 6th International Conference on Web Information Systems and Technologies (WEBIST), Valencia, Spain. pp. 13–20 (2010).
15. Meroño-Peñuela, A., Wittek, P., Darányi, S.: Visualizing the Drift of Linked Open Data Using Self-Organizing Maps. In: Drift-a-LOD Workshop at the 20th International Conference on Knowledge Engineering and Knowledge Management (2016).
16. Guarino, N., Welty, C.: A Formal Ontology of Properties. In: Knowledge Engineering and Knowledge Management Methods, Models, and Tools. pp. 97–112. Springer Berlin Heidelberg (2000).
17. Gangemi, A., Guarino, N., Masolo, C., Oltramari, A., Schneider, L.: Sweetening ontologies with DOLCE. In: International Conference on Knowledge Engineering and Knowledge Management. pp. 166–181. Springer Berlin Heidelberg (2002).
18. Monge, A.E., Elkan, C.: The Field Matching Problem: Algorithms and Applications. In: 2nd Intl. Conf. Knowledge Discovery and Data Mining (KDD). pp. 267–270 (1996).
19. Ashkpour, A., Meroño-Peñuela, A., Mandemakers, K.: The aggregate Dutch historical censuses: Harmonization and RDF. *Hist. Methods A J. Quant. Interdiscip. Hist.* 48, 230–245 (2015).
20. Meroño-Peñuela, A., Ashkpour, A., Guéret, C., Schlobach, S.: CEDAR: the Dutch historical censuses as linked open data. *Semant. Web.* 1–14 (2015).
21. Miles, A., Bechhofer, S.: SKOS simple knowledge organization system refer-ence. (2009).
22. Leeuwen, M.H.D. van, Maas, I., Miles, A.: HISCO: Historical international standard classification of occupations. Leuven: Leuven University Press (2002).
23. Meroño-Peñuela, A., Guéret, C., Hoekstra, R., Schlobach, S.: Detecting and reporting extensional concept drift in statistical linked data. Presented at the (2013).