Quality measures for `skos:exactMatch` linksets: an application to the thesaurus framework LusTRE

RICCARDO ALBERTONI

Istituto di Matematica Applicata e Tecnologie Informatiche,
Consiglio Nazionale delle Ricerche,
Via De Marini, 6, Genova, 16149, Italy, riccardo.albertoni@ge.imati.cnr.it

MONICA DE MARTINO

Istituto di Matematica Applicata e Tecnologie Informatiche
Consiglio Nazionale delle Ricerche,
Via De Marini, 6, 16149 Genova, Italy, demartino@ge.imati.cnr.it

PAOLA PODESTÀ

Istituto di Matematica Applicata e Tecnologie Informatiche,
Consiglio Nazionale delle Ricerche,
Via De Marini, 6, 16149 Genova, Italy, podesta@ge.imati.cnr.it

Received (04 05 2017)
Revised (28 02 2018)

Purpose- The paper focuses on the quality of the connections (linkset) among thesauri published as Linked Data on the Web. It extends the cross-walking measures with two new measures to evaluate the enrichment brought by the information reached through the linkset (lexical enrichment, browsing space enrichment). It fosters the adoption of cross-walking linkset quality measures besides the well-known and deployed cardinality-based measures (linkset cardinality, linkset coverage).
Design/Methodology/Approach- The paper applies the linkset measures to the Linked Thesaurus fRamework for Environment (LusTRE). LusTRE is selected as testbed as it is encoded using Simple Knowledge Organization System (SKOS) published as Linked Data, and it explicitly exploits the cross-walking on its validated linksets.
Findings- The application on LusTRE offers an insight of the complementarities among the considered linkset measures. In particular, it shows that the cross-walking measures deepen the cardinality-based measures analysing quality facets that were not previously considered. The actual value of LusTRE's linksets regarding the improvement of multilingualism and concept spaces is assessed.
Research limitations/implications- The paper considers `skos:exactMatch` linksets, which belong to a rather specific but a quite common kind of linkset. The cross-walking measures explicitly assume correctness and completeness of linksets. Third party approaches and tools can help to meet the above assumptions.
Originality/value- This paper fulfils an identified need to study the quality of linksets. Several approaches evaluate Linked Data quality focusing on dataset quality but disregarding the other essential component: the connections among data.

Keywords: Quality; Linkset; SKOS; Linked Data; Environmental Thesauri; Cross-walking.

Research paper

1. Introduction

In the paper "Linked Data - The Story So Far", Bizer et al. (2009) were among the first to take a picture of the enormous transformation of the Web of Document into the Web of Data. Since then, the Linked Data popularity has never ceased to grow. Linked Data aims at disclosing the potential of independently served data dealing with access and integration issues. It publishes documents encoded using Resource Description Framework (RDF), but also "uses RDF to make typed statements that link arbitrary things in the world. The result, which we will refer to as the Web of Data, may more accurately be described as a web of things in the world, described by data on the Web" (Bizer et al., 2009). Linked Data allows RDF data to be published, shared, retrieved, reused and analysed unlocking the existing data silos to a broader community of consumers. RDF provides a graph-based data model based on triples in the form of *subject, predicate, object* (Schreiber and Raimond, 2014). Both data and links among data are expressed with triples. Linked Data relies on two fundamental Web technologies: the Internationalized Resource Identifiers (IRIs[1]) and the HyperText Transfer Protocol (HTTP), which are respectively deployed as the global identifiers and as the protocol to dereference the information that is associated with each identifier. Following Linked Data principles[2] several billions of facts encoded in RDF triples have been published in the Linked Open Data (LOD) Cloud[3].

This vast quantity of newly available and connected datasets is transforming the Web into a global data space enabling new types of analysis and applications in diverse domains including Life Science, Government, Environment, and Cultural Heritage. At the same time, the evaluation of the quality of these newly served data becomes critical. "Data quality can affect the potentiality of the applications that use data. As a consequence, its inclusion in the data publishing and consumption pipelines is of primary importance" (Calegari et al., 2017). The challenge is twofold: to evaluate the quality of the data on the Web and to make quality-related information explicit, understandable and consumable to both humans and machines.

Several existing initiatives have the goal to define new metrics and to evaluate the quality of Linked Data. The W3C Data Quality Vocabulary (DQV) (Albertoni and Isaac, 2016) introduces a common way to document the quality of a dataset, making easier to publish, exchange and consume quality metadata. Recent works such as Zaveri et al. (2016); Debattista et al. (2016*b*); Radulovic et al. (2018); Kontokostas et al. (2014) consider different aspects of Linked Data quality, called dimensions, e.g., accessibility, interlinking, performance, syntactic validity or completeness. They define and deploy several concrete metrics (or measures) to precisely and objectively evaluate each dimension. However, they focus on Linked Data datasets, reserving very limited attention to their connections, the *linksets*. A linkset is a set of homogeneous links, all of the same types and connecting the same *subject dataset* to the same *object dataset* (Alexander et al., 2011). The quality of linksets is studied as part of the *interlinking dimension* defined in the recent state of the art (Zaveri et al., 2016). Few metrics are defined to evaluate interlinking, they mainly focus on correctness (e.g., broken links, open `owl:sameAs` chains, crowdsourcing method), or on the number of links (linkset cardinality), or on the extent to which a linkset covers the elements of a dataset (linkset coverage) (Guéret et al., 2012; Zaveri et al., 2016; Albertoni and Gómez Pérez, 2013).

The experience gained creating LusTRE[4], the multilingual linked thesaurus framework for the environment, has taught us to pay attention to the quality of connections between datasets. LusTRE has been designed during the EU project eENVplus[5] extending and redesigning the Common Thesaurus Framework for the Environment (De Martino and Albertoni, 2011). LusTRE faces with cross-lingual and cross-sectoral issues in environmental data sharing: it provides a wide multilingual terminology obtained by linking available thesauri for the different disciplines in the Environment and a set of web services to exploit them (Albertoni et al., 2018). The eENVplus project has spent considerable efforts to review the available environmental thesauri checking those not yet available as linked data (Albertoni et al., 2014*a*). Then, it has published ThiST[6] and EARTh (Albertoni et al., 2014*b*) as Linked Data using the Simple Knowledge Organization System (SKOS) (Miles and Bechhofer, 2009), and connected them to GEMET[7], AGROVOC (Caracciolo et al., 2013) and EUROVOC[8].

In LusTRE, the linksets among the thesauri are particularly important as they are exploited to satisfy user requests. LusTRE enriches user navigations and service results with translations and concepts which are reachable through the linkset. Thus, the linkset quality becomes a critical issue. Given a linkset between two SKOS thesauri, LusTRE should evaluate the multilingual enrichment obtained in terms of newly translated labels reachable through a linkset. This information helps to address the incomplete language coverage issue, which affects many popular SKOS thesauri (Suominen and Mader, 2014). It also needs to evaluate the number of new concepts reached by crossing a linkset, as this helps to assess the enrichment of the space of concepts that can be browsed (aka, the thesaurus browsing space). The existing quality measures do not address the above needs: they mainly focus on interlinking among datasets made by `owl:sameAs` relations, and they do not consider the quantity and the kind of new information reachable through the linkset.

Motivated by these limitations, this paper presents a set of measures to evaluate the quality of linksets made only by `skos:exactMatch` links, aka `skos:exactMatch` linksets. The `skos:exactMatch` is the most-commonly used link type in LusTRE. It relates concepts that can be used interchangeably across a wide range of information retrieval application.

The contributions of the paper are:

- It presents new linkset quality measures: the *lexical enrichment* and the *browsing space enrichment*. Such measures calculate the average number of new lexical values or of new concepts reachable in the object thesaurus per concept of the subject thesaurus. The paper considers such kind of information as the impact on the subject thesaurus of the new values reached through the linkset. These two measures extend the *average linkset importing* and the *average linkset reachability* presented in our earlier work (Albertoni et al., 2016), which capture the quantity of new information reached crossing the linkset but neglect the impact on the subject thesaurus. These four measures are classified as *cross-walking measures*

since they are based on the quantity and kind of information reached navigating the linkset. They require the correctness of links and also the linkset completeness. Linkset completeness means that any concept in the subject thesaurus having an exact equivalent concept in the object thesaurus must be involved in a `skos:exactMatch` link in the linkset.
- It revises two measures available in literature used to evaluate the linksets in Linked Data (linkset cardinality, linkset coverage), and it identifies such measures as *cardinality-based measures* since they are based on the number of links composing the linkset.
- It applies all the considered linkset quality measures to the linksets of the thesaurus framework LusTRE bringing a two-fold contribution. On the one hand, it compares the measures discussing their relations. It highlights that *cardinality-based measures* often give a very "poor" and sometimes misleading characterization of the quality justifying the need for more accurate investigations by applying the *cross-walking measures*. On the other hand, the results of the *cross-walking measures* highlight the values of LusTRE's linksets. They measure the amount of new information reachable by each linkset as well as the impact that this information has on the subject thesaurus.

The paper is organised as follows. Section 2 presents the related work on Linked Data quality. Section 3 illustrates all the measures considered for linkset quality. Section 4 describes the framework LusTRE. Section 5 discusses the results obtained applying the measures to the LusTRE's linksets. It discusses the differences and mutual relations among the measures, and it evaluates the LusTRE's linksets assessing their multilingual and browsing space improvement. Section 6 presents conclusions and future work.

2. Related Work

Data quality assessment is a multidimensional issue affecting data integration, sharing, retrieval and analysis, in fact, as mentioned by Pipino et al. (2002), "if stakeholders assess the quality of data as poor, their behaviour will be influenced by this assessment". Al-Hakim (2006) presents the common notion of data quality as the "fitness for intended use in operations, decision-making, and planning". It highlights that the quality assessment depends on general factors, aka dimensions, e.g., accessibility, security, timeliness, consistency, and, on task-dependent factors related to the consumers' satisfaction.

Recent initiatives, such as the Data on the Web Best Practices W3C working group, have pointed out that not only the evaluation of web data quality but also the publication, exchange and consumption of quality information are crucial (Calegari et al., 2017). To this aim, the group has developed the Data Quality Vocabulary (DQV) (Albertoni and Isaac, 2016) which is a specific RDF vocabulary to document the quality of data on the web. The sections below present works specifically related to Linked Data and SKOS.

2.1. *Linked Data quality assessment*

The need for Linked Data specific quality assessment methods arises from the amazing quantity of data published and exchanged on the Web using the Linked Data paradigm in the last few years (several billion of triples in the LODCloud[9]). Zaveri et al. (2016) present a comprehensive review of the approaches that focus on Linked Data, becoming the landmark of several following papers. Zaveri et al. (2016) borrow the terminology introduced in Bizer and Cyganiak (2009). They identify, with *dimensions*, the characteristics of datasets that are relevant to consumers, and with *indicators* (measure or metric), the procedures for measuring a quality dimension. They analyse 21 existing approaches published from 2002 to 2014 extracting 18 different quality dimensions and their definition, revising several metrics for each dimension (69 metrics in total). Radulovic et al. (2018) present a unified Linked Data quality model according to the terminology described in the ISO standards (ISO/IEC 15939:2007 and ISO/IEC 25010). They extend the work of Zaveri et al. (2016) presenting a larger set of quality characteristics and measures. They model three different measurement levels (triple, graph, dataset) and also highlight some interesting relationships between the measures at different levels. They also extend the DQV to capture quality information specific to Linked Data.

Both the previous works focus on datasets, while linksets are not fully investigated. On the one hand, Radulovic et al. (2018) do not consider interlinking as a dataset quality dimension, but they use the measures for interlinking presented in Zaveri et al. (2016) to evaluate the *completeness* dimension. On the other hand, Zaveri et al. (2016) define the quality dimension *interlinking* as "the degree to which entities that represent the same concept are linked to each other, be it within or between two or more data sources". Moreover, it is worth to note that the notions of interlink and linkset are closely related but not the same. A linkset is a set of links which have the same type, from the same subject dataset to the same object dataset. An interlink is a set of links without any assumption about the link's types, where each link might in principle involve distinct datasets.

Guéret et al. (2012) present one of the early works on the assessment of the *interlinking* dimension. They present three classic network measures (i.e., degree, centrality, clustering coefficient) and two measures specifically designed for Linked Data (i.e., *Open SameAs* chains, and *Description Richness*). These measures are implemented in the framework LINK-QA for determining whether a set of `owl:sameAs` links improves the overall quality of Linked Data. Among the measures, *Open SameAs* chains are useful for assessing the link correctness while *Description Richness* simply counts the new RDF triples reached through `owl:sameAs` links. The authors admit that classic networks measures can only partially detect the quality of links, as such measures rely only on the topology of the network without considering the semantics of links. In particular, these measures evaluate `owl:sameAs` links without distinguishing the kind of information reached (e.g., if the triples add lexical terms or entities of specific types). Instead, Albertoni and Gómez Pérez (2013) focus on linkset quality considering the dimension *completeness*. They present novel measures: *Linkset Type Coverage* and *Linkset Entity Coverage for Type* evaluating the coverage of the `owl:sameAs` linkset on different information in a dataset. However, SKOS links such as `skos:exactMatch` are not explicitly considered.

Part of the literature on Linked Data quality presents automated tools implementing the existing measures and the applications of such measures to set of Linked Data datasets. Kontokostas et al. (2014) present RDFUnit, which is a framework that supports a quality assessment based on test cases encoded as SPARQL queries (Harris and Seaborne, 2013). However, the quality of interlinking is not specifically addressed in this framework. Two crowdsourcing-based methodologies and a tool for dataset quality assessment are presented in Acosta et al. (2013), and Zaveri et al. (2013) respectively, and then, applied to DBpedia. These two works evaluate interlinking only considering the correctness of links. The framework Luzzu, which is presented by Debattista et al. (2016*a*), implements the metrics reviewed by Zaveri et al. (2016) and provides the possibility to add domain specific quality metrics. It includes a set of comprehensive ontologies for capturing and exchanging data quality information, and a user-driven quality based weighted ranking algorithm. Finally, Färber et al. (2018) extend the quality dimensions proposed in Wang and Strong (1996), with *consistency* and *verifiability*, while the *accessibility* dimension is improved with the Linked Data specific measures *license* and *interlinking* defined in Zaveri et al. (2016). Thus, according to the 34 data quality criteria existing in the literature, they analyse the datasets DBpedia, Freebase, OpenCyc, Wikidata, and YAGO proposing a framework which enables users to find the most suitable dataset for their needs. They measure the interlinking on instance level by calculating the extent to which instances have at least one `owl:sameAs` link to external resources, a notion that is similar to the coverage previously discussed. They also check for the dereferenceability of links' objects.

2.2. *SKOS quality assessment*

Several efforts are undertaken to specifically evaluate the quality of the Knowledge Organization Systems (KOSs), such as thesauri. KOSs are very often used to share standard and scientific/technical terms, also in multiple languages to improve discoverability and semantic interoperability of data among the different communities. For example, AGROVOC (Caracciolo et al., 2013) is a thesaurus covering all areas of interest of the Food and Agriculture Organization (FAO) of the United Nations; EARTh (Albertoni et al., 2014*b*) is a general-purpose thesaurus for the environment; GACS Core (Baker et al., 2016) provides a concept scheme, where their most frequently used concepts from agricultural thesauri are interlinked and integrated.

Lacasta et al. (2016) present a process to automatically analyse the quality of thesauri, considering in particular Urban and GEMET, according to ISO 25964. They perform lexical and syntactic analysis of the lexical properties (preferred label and description) and an analysis of the broader/narrower relations. Albertoni et al. (2014*a*) present a methodological approach based on the 5 Star Linked Data principles to evaluate the reusability of thesauri. Both the previous works do not consider any quality measures for linksets.

Suominen and Hyvönen (2012) present several quality issues affecting SKOS thesauri, and they also create and use the Skosify[10] tool to improve the validity of fourteen SKOS vocabularies automatically correcting the majority of structural problems. Suominen and

Mader (2014) present 26 quality issues automatically computable together with a set of heuristics to automatically correct 12 of such issues. They present how to improve the quality of thesauri, overcoming these issues with quality assessment tools such as qSKOS (Mader et al., 2012), the PoolParty[11] checker and the quality improvement tool Skosify. Quarati et al. (2017) propose a methodology to support the context driven quality assessment of the aforementioned issues. The only metrics assessing the interlinking quality are the out-links and in-links, which do not evaluate the new lexical labels or the new concepts reached through the linkset. Our previous work (Albertoni et al., 2015) presents a first attempt to evaluate the `skos:exactMatch` linkset quality between two SKOS thesauri considering lexical information. The paper focuses on the multilingual gain in terms of newly translated labels obtained by complementing a SKOS thesaurus through `skos:exactMatch` links. To this purpose, the *linkset importing* measure is presented and applied to the EARTh's linksets. In our following work (Albertoni et al., 2016), the linkset importing has been refined with two new linkset measures: the *average linkset importing* and the *average linkset reachability*, which are reused and extended in this paper. The first measure evaluates the multilingual enrichment obtained by newly translated labels reachable in the object thesaurus through a linkset. It can help to overcome the incomplete language coverage issue, that is, when `skos:prefLabel` and `skos:altLabel` are provided in all the expected languages only for a subset of the thesaurus concept (Suominen and Hyvönen, 2012). The second one evaluates the number of new concepts reached in the object thesaurus by crossing a linkset. It can be exploited to evaluate the potential enrichment of the space of browsable concepts. These average measures are the first evaluating the quantity and the kind of new information reached through the linkset, but they do not evaluate the quantity of new information available for each concept in the subject thesaurus.

3. Linkset Quality Measures

This section formalises the linkset quality measures deployed in the rest of this paper. It recalls the literature definitions for cardinality-based measures: *linkset coverage* and *linkset cardinality*. It introduces four cross-walking measures: the *average importing* and the *average reachability* presented in our recent work (Albertoni et al., 2016), and the *lexical enrichment* and the *browsing space enrichment* which are two novel measures complementing the previous.

The measures assess the quality of linksets between thesauri encoded using SKOS. SKOS is a W3C RDF/OWL vocabulary, which encodes concepts, concepts' lexical representation and relations. Concepts are identified using URIs and labelled with multilingual lexical labels such as preferred labels (i.e., `skos:prefLabel`), and alternative labels (i.e., `skos:altLabel`). Concepts are related to other concepts, for example, they can be organised into hierarchies using the properties `skos:broader` and `skos:narrower` or into association networks using the property `skos:related`. SKOS allows expressing different kinds of links among concepts belonging to independently defined thesauri: the property `skos:exactMatch` links identical concepts,

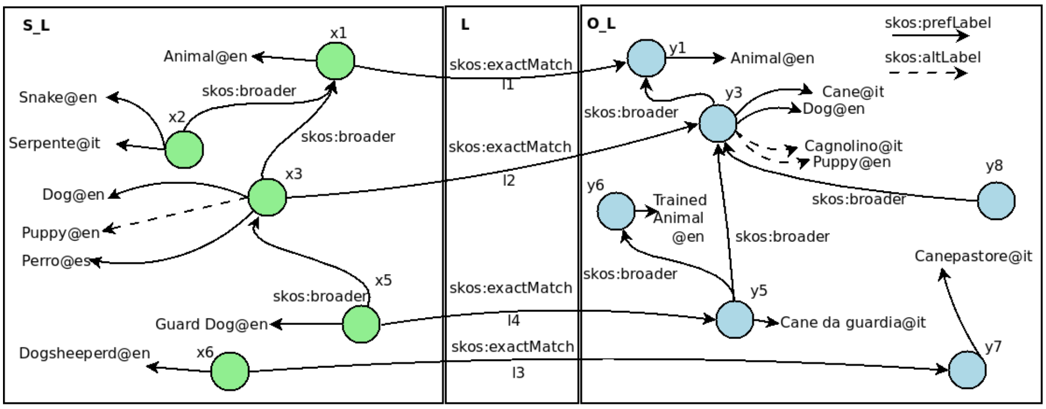

Fig. 1. An example of RDF/SKOS thesauri and `skos:exactMatch` linkset.

the property `skos:closeMatch` links concepts that almost identical, the properties `skos:broadMatch`, `skos:narrowMatch`, `skos:relatedMatch` specify other kinds of matching between concepts.

The paper relies on the notion of linkset provided in the Vocabulary of Interlinked Datasets (VoID) (Alexander et al., 2011), in which a **linkset (L)** is a special kind of dataset containing only RDF links between the `void:subjectsTarget` and the `void:objectsTarget` respectively representing the **object** and the **subject** of the linkset. Each RDF link is an RDF triple *(s,p,o)*, where *s* and *o* belong respectively to the *subject* and *object* dataset of the linkset, *p* is an RDF property that indicates the type of the link. RDF links of a linkset should all have the same type, otherwise, the linkset should be split into distinct linksets. In particular, this paper considers ***skos:exactMatch* linksets**, which are linksets made only by `skos:exactMatch` links among **SKOS thesauri**. Table 1 grouds some of the basic definitions needed for defining our metrics.

The following sections introduce the *cardinality-based measures* and the *cross-walking measures*.

3.1. *Cardinality based measures*

The cardinality-based measures evaluate the linkset quality through the number of links included in the linkset. This work considers the following cardinality-based measures:

- the *linkset cardinality*, which counts the number of `skos:exactMatch` links in the linkset. It is identified with $|L|$;
- the *linkset coverage*, which returns the proportion of concepts in the subject thesaurus involved in the linkset. It is formally defined as $Cov_L(S_L, L) = \frac{|L|}{|C(S_L)|}$.

3.2. *Cross-walking measures*

The cross-walking measures evaluate the linkset quality considering the quantity and kind of new information reached in the object thesaurus navigating the linkset. Firstly, this paper

Table 1. Basic formal definitions

Symbols	*Definition*
$(s, p, o) \in$ (RDFRiri$\cup$BNode)$\times$RDFProp$\times$(RDFRiri$\cup$ BNode $\cup$ RDFLit)	an RDF triple where s, p, o are RDF terms, respectively the subject, the property and the object. RDFRiri is the set of RDF resources denoted by an IRI, BNode is the set of blank nodes, and RDFLit is the set of RDF literals. RDFProp represents the set of RDF properties
S_L	the linkset subject thesaurus, which is the set of triples contained in the subject thesaurus of the linkset L
S_l	the subject concept of a link l
O_L	the linkset object thesaurus, which is the set of triples contained in the object thesaurus of the linkset L
O_l	the object concept of a link l
$C(T)$	the SKOS concepts defined in a thesaurus T
$SKOSRel$ ={skos:broader, skos:related, skos:narrower}	the set of SKOS properties connecting SKOS concepts in a thesaurus
$SKOSLabel$ ={skos:prefLabel, skos:altLabel, skos:hiddenLabel}	the set of RDF properties connecting concepts to lexical labels
L ={$(s$, skos:exactMatch, $o)\|s \in C(S_L), o \in C(O_L)$}	the linkset, which is the set of triples connecting a subject thesaurus S_L to object thesaurus O_L by `skos:exactMatch`
L_O={$o \in O_L\|(s$, skos:exactMatch, $o) \in L$}	the set of objects in the linkset
$val4P_T(c, p, ln)$	the set of lexical labels in a thesaurus T which are in a language ln and are associated to a concept c by a $SKOSLabel$ property p. The language parameter ln is specified by using the standard ISO language tags or the symbol _ to indicate the whole set of ISO languages. **Examples:** Considering Figure 1, for the link l_2, val4P$_{O_{l2}}$(y$_3$,`skos:prefLabel`,en)={Dog@en}, whilst val4P$_{S_{l2}}$(x$_3$,`skos:prefLabel`,_) ={Dog@en, Perro@es}, since in the latter there is no constraint on the language tag.
RC(k,SR,L_O,O_L)	the set of concepts in the object thesaurus (O_L) reachable by the concepts in L_O, through the relations in $SR \subseteq SKOSRel$, in some hops $\leq$ k. **Example:** Considering Figure 1, a number of hop k equal to 2 and SR={`skos:broader`}, $RC(2$,{`skos:broader`}$, L_o, O_L)$= {y_1, y_3, y_5, y_6, y_7}.

refers to *average linkset importing* and the *average linkset reachability* (Albertoni et al., 2016) to evaluate the average number of new information accessed navigating each link of the linkset. However, these measures do not evaluate how such as information enrich the subject thesaurus. Thus, this paper defines two new measures: the *lexical enrichment* and the *browsing space enrichment*, which evaluate how significant is the impact of the new information on the subject thesaurus. For example, "one new `skos:prefLabel` in English for each concept in the subject thesaurus" should be considered more impacting than "one new `skos:prefLabel` in English every ten concepts". The cross-walking measures are deployed to evaluate the improvements brought by traversing a linkset. In particular, the average linkset importing and lexical enrichment evaluate the multilingual gain, that is, the number of newly translated labels reached and the impact that they have on the subject. While the average linkset reachability and the browsing space enrichment evaluate the number of new concepts reached and how these concepts widen the set of concepts already available in the subject thesaurus.

All these cross-walking measures require the correctness and completeness of linksets. Otherwise, these measures might consider duplicated information which leads to imprecise results. Correctness and completeness are reasonable assumptions: (i) currently, all applications consuming Linked Data implicitly assumes at least correctness (trusting on publisher reliability); (ii) tools like SILK[12] and LIMES[13] help to reach these assumptions.

Definition 1. ***(Average linkset importing).*** *Let ln be a particular ISO language tag or _ for indicating all ISO languages, and* $p \in SKOSLabel$*, l a link in L of the form (*c_{S_l}*, skos:exactMatch,* c_{O_l}*). The average linkset importing is defined as follows:*

$$ALI_L(p, ln) = \frac{1}{|L|} \sum_{l \in L} |val4P_{O_l}(c_{O_l}, p, ln) \setminus val4P_{S_l}(c_{S_l}, p, ln)|$$

It evaluates the average number of new preferred labels or alternative labels (depending on which $SKOSLabel$ *property is indicated as p) brought by each link traversed in the linkset L.*

Example 3.1. *Considering the thesauri* S_L*,* O_L*, the linkset* L *in Figure 1, the links* l_2*,* l_3*,* l_4 *import one new Italian* `skos:prefLabel` *each (respectively, "Cane@it", "Cane Pastore@it", "Cane da Guardia@it"), whilst* l_1 *does not import any label. In fact,* ALI_L*(*`skos:prefLabel`*, it)=* $\frac{1}{4} * (0+1+1+1)$*=* $\frac{3}{4}$*=0.75, which means that three new Italian translations for* `skos:prefLabel` *are reachable every four links.*
Disregarding the language of the imported preferred labels results in the same, ALI_L*(*`skos:prefLabel`*, _)=0.75 : the aforementioned Italian* `skos:prefLabel` *are the only that can be imported, as all the English* `skos:prefLabel` *which could be additionally imported from the links' objects (i.e., "Dog@en" "Animal@en") are already in the links subjects.*

Definition 2. ***(Average linkset reachability).*** *Let* $SR \subset SKOSRel$ *be the set of SKOS relations considered relevant, and k the number of hops. The average linkset reachability is defined as follows:*

$$ALR_L(SR,k) = \frac{1}{|L|} * |RC(k, SR, L_O, O_L) \setminus L_O|$$

It evaluates the average number of concepts reachable by cross-walking a link of the linkset and exploring the object thesaurus O_L *until a certain depth, identified with the number of hops. Concepts in the set* L_O *are not counted.*

Example 3.2. *Considering the thesauri* S_L *and* O_L*, and the linkset L in Figure 1,* y_6 *is the only new node reached via* `skos:broader` *properties in two hops:* y_8 *cannot be reached* *by the* `skos:broader`*, it could have been reached by its inverse, the* `skos:narrower`*, but the inverse is not included in the considered properties, and all the remaining nodes are equivalent to some node in the subject thesaurus as they are part of* L_O*=*{y_1, y_3, y_5, y_7}*. The* $ALR(SKOSrel, 2)$= $\frac{1}{4} * |\{y_1, y_3, y_5, y_6, y_7\} \setminus \{y_1, y_3, y_5, y_7\}| = \frac{1}{4} * |\{y_6\}| = \frac{1}{4}$*=0.25. The 0.25 represents the average number per link of new concepts reachable, in 2 hops, in the object thesaurus. This result means that one new concept is reachable every four links.*

Definition 3. ***(Lexical terms enrichment).*** *Let ln be an ISO language tag or _, and* $p \in SKOSLabel$*, l a link in L of the form (*c_{S_l}*, skos:exactMatch,* c_{O_l}*). The lexical enrichment for considering property p and language ln, through L, is defined as follows:*

$$LTE_L(p, ln) = \frac{1}{|C(S_L)|} \sum_{l \in L} |val4P_{O_l}(c_{O_l},p,ln) \setminus val4P_{S_l}(c_{S_l},p,ln)|$$

It evaluates the number of new lexical values reached in the object thesaurus with respect to the total concepts in the subject thesaurus, that is the average number of new values for each concept in the subject thesaurus.

Example 3.3. *Considering Figure 1, the new* `skos:preflabel` *in Italian reachable in* O_L *through the linkset are {cane@it, canepastore@it, canedaguardia@it}. While, the number of concepts in the subject* $|C(S_L)|$ *is 5. Thus,* LTE_L*(*`skos:preflabel`*,it) =* $\frac{3}{5}$ = 0.6 *means three new* `skos:preflabel` *in Italian every five concepts in the subject.*

Definition 4. ***(Browsing space enrichment).*** *Let* $SR \subset SKOSRel$ *be the set of SKOS relations considered relevant, and k the number of hops. The browsing space enrichment is defined as follows:*

$$BSE_L(SR,k) = \frac{1}{|C(S_L)|} * |RC(k, SR, L_O, O_L) \setminus L_O|$$

It evaluates the number of new concepts reached in the object thesaurus in k *hops compared to the total number of concepts in the subject thesaurus.*

Example 3.4. *Considering Figure 1,* y_6 *is the only node reached in 2 hops with* `skos:broader` *properties, but the cardinality of concepts in* S_L *is* $C(S_L)$ *= 5. Thus,* BSE_L*(*`skos:broader`*,2) =* $\frac{1}{5} * |\{y_1, y_3, y_5, y_6, y_7\} \setminus \{y_1, y_3, y_5, y_7\}| = \frac{1}{5}$*=0.2. The 0.2 represents the average number of new concepts available for each concept in the subject thesaurus, that is one new concept every five concepts.*

Two remarks help a deeper understanding of the measures: (i) the *linkset coverage* is the multiplier which transforms the *average linkset importing* into the *lexical term enrichment* and the *average linkset reachability* into the *browsing space enrichment*. Although this relation holds at algebraic level, it is worth to be noted to comprehend the nature and the purpose of the measures. The multiplier switches the comparison base moving the assessment from the number of new imported labels /new reached concepts per link to the actual impact they have on the linkset subject. (ii) the results of the cross-walking measures differ whether applied to a linkset or its inverse (i.e., the linkset obtained inverting the direction of the `skos:exactMatch`). That is due to links as l_2 in Figure 1: l_2 imports the English `skos:altLabel` "cagnolino@it" from y_3 whilst the inverse link of l_2 does not import any `skos:altLabel`. The quality of a linkset and the quality of its inverse are unrelated for the cross-walking measures.

The cross-walking measures have been implemented in a research prototype developed in JAVA. The prototype provides a command-line tool to assess the linkset quality, which exploits the JENA framework to connect either to thesauri's and linksets' dumps or to their SPARQL endpoint.

4. LusTRE: a case of study

The Linked Thesaurus fRamework for the Environment (LusTRE) (Albertoni et al., 2018) is an interesting case study for the exploitation of Linked Data to support metadata compilation and information discovery. The framework is implemented in compliance with the Linked Data principles. It provides a knowledge infrastructure of interlinked thesauri for the Environment. The following components characterise LusTRE: (i) **LusTRE Knowledge Infrastructure** contains different environmental SKOS vocabularies and the linksets among them, enabling the access as one virtual integrated Linked Data source. It is deployed on a Virtuoso server and accessible by SPARQL endpoint. (ii) **LusTRE Exploitation Services** is a set of end-user oriented web services with a REST interface. These services allow exploiting the knowledge contained in the LusTRE knowledge infrastructure for improving client applications such as a metadata editor or a geodata portal. For example, EUOSME[14], QSPHERE[15] and the under development version of the INSPIRE Geoportal[16] integrate the LusTRE web services. (iii) **LusTRE Web Interface** provides a human-accessible interface to manually search and navigate the interlinked knowledge infrastructure using textual or visual browsing. In particular, LusTRE Exploitation Services and Web Interface support the *cross-walking* among vocabularies. The *cross-walking* allows automatically to navigate among matching concepts belonging to different linked thesauri. It works beyond the scope and limitations of a single thesaurus, possibly enriching data at hand, and, thus, improving user satisfaction in data consuming process. LusTRE includes the following SKOS thesauri: ThIST (34150 concepts, 2 languages), AGROVOC (32310 concepts, 24 languages), EARTh (14350 concepts, 2 languages), EUROVOC (6883 concepts, 23 languages), GEMET (5223 concepts, 32 languages). Concerning linksets, it provides the twenty `skos:exactMatch` linksets presented in Table 2. The table describes the pairs of thesauri involved and the cardinality of each linkset. Linksets are created by working out the transitive closure on existing `skos:exactMatch` and applying specific linkset discovery tasks to ensure linkset correctness. Every linkset discovery task follows a two-steps process: firstly, SILK discovers new links, and then the link correctness is validated by domain experts. In particular, for the linksets between ThIST and AGROVOC, ThIST and EUROVOC, ThIST and GEMET the experts have accepted in average the 99% of links discovered by SILK. Linkset completeness is reasonably ensured by having applied different and not very restrictive matching functions during the discovery

Table 2. Linksets in LusTRE.

Linkset Name	*Subject Thesaurus*	*Object Thesaurus*	*Linkset Cardinality*
AGROVOCTOEARTh	AGROVOC	EARTh	1438
AGROVOCTOEUROVOC		EUROVOC	1269
AGROVOCTOGEMET		GEMET	1188
AGROVOCTOThIST		ThIST	1695
EARThTOAGROVOC	EARTh	AGROVOC	1438
EARThTOEUROVOC		EUROVOC	1346
EARThTOGEMET		GEMET	4365
EARThTOThIST		ThIST	1140
EUROVOCTOAGROVOC	EUROVOC	AGROVOC	1269
EUROVOCTOEARTh		EARTh	1346
EUROVOCTOGEMET		GEMET	1683
EUROVOCTOThIST		ThIST	733
GEMETTOAGROVOC	GEMET	AGROVOC	1188
GEMETTOEARTh		EARTh	4365
GEMETTOEUROVOC		EUROVOC	1683
GEMETTOThIST		ThIST	792
ThISTTOAGROVOC	ThIST	AGROVOC	1695
ThISTTOEARTh		EARTh	1140
ThISTTOEUROVOC		EUROVOC	733
ThISTTOGEMET		GEMET	792

task. SKOS entailments have been materialised to support clients with limited processing power. As a consequence of such materialisation and the `skos:exactMatch` symmetry, reciprocal linksets (e.g., EARThTOGEMET and GEMETTOEARTh or EUROVOCTOAGROVOC and AGROVOCTOEUROVOC in Table 2) have the same links but inverted. Concerning the quality of the subject and object thesauri, LusTRE directly involved the developer of two thesauri (ThIST and EARTh), and it trusted in the producers of other thesauri (e.g., AGROVOC, EUROVOC, GEMET). This choice has also been made considering that the focus of this paper is not the quality of datasets but the quality of linksets. Moreover, in other contexts such as the LOD Cloud where the trust to producers cannot be granted, the existing tools (e.g., qSKOS) can evaluate the quality of datasets.

5. Application and discussion

The measures described in Section 3 were applied to assess the quality of `skos:exactMatch` linksets in LusTRE. The average linkset importing and lexical enrichment were applied considering `skos:prefLabel` and `skos:altLabel` in different languages. While, the average linkset reachability and the browsing space enrichment were applied navigating at most 4 hops traversing the SKOS relations `skos:narrower`, `skos:broader` and `skos:related`. The application considered 4 hops (i.e., k=4) since it was indicated by the eENVplus stakeholders as a reasonable number of steps to search for the specialisation/generalisation of a particular SKOS term in a thesaurus.

This section illustrates the quality results according to two perspectives: the properties of the quality measures and the quality of LusTRE's linksets. The first perspective dis-

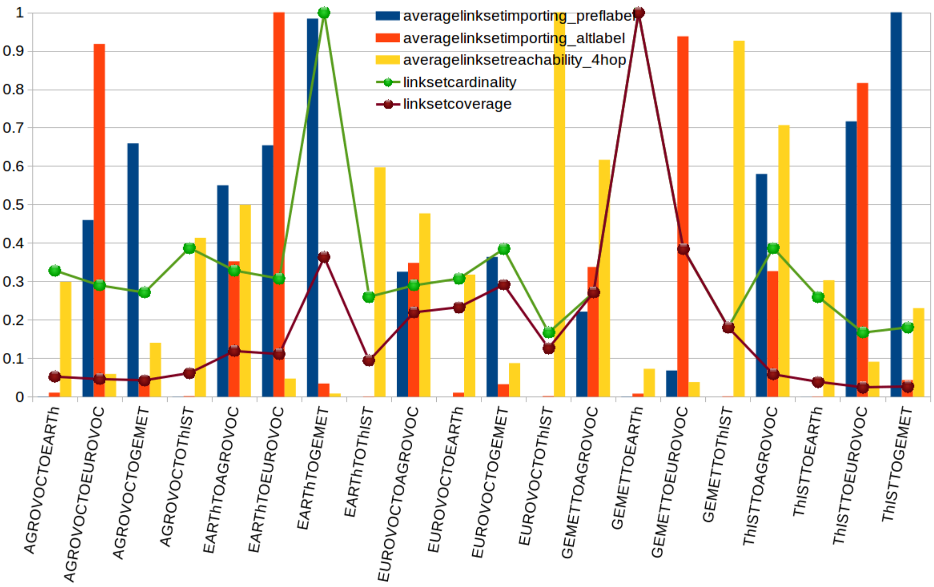

Fig. 2. Results of the average linkset reachability and the average linkset importing for `skos:prefLabel` and `skos:altlabel` compared with cardinality-based measures applied to LusTRE. The results are normalised between 0 and 1.

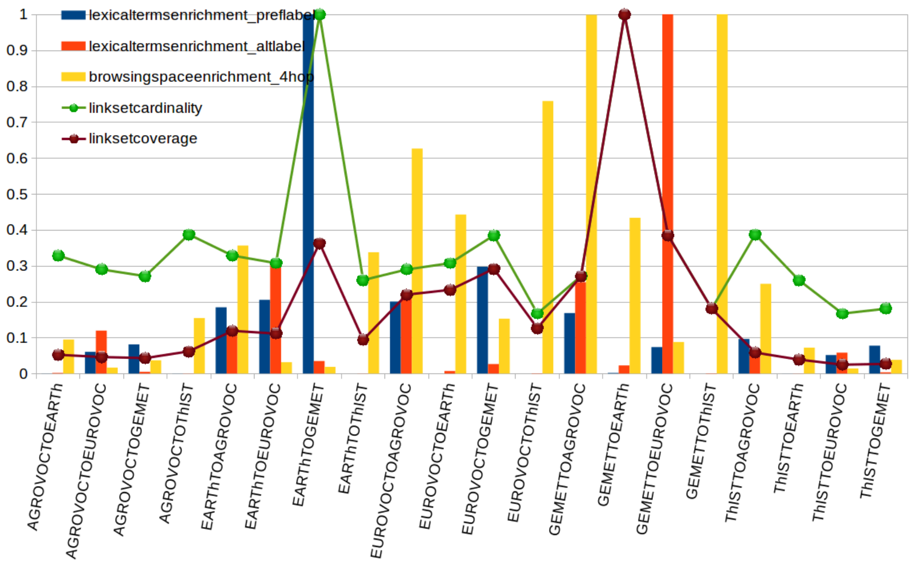

Fig. 3. Results of the browsing space and lexical enrichment compared with cardinality-based measures applied to LusTRE. The results are normalised between 0 and 1.

cusses the relations and the differences among the quality measures. The second perspective presents the multilingual and browsing space's improvement brought by the linksets to the thesauri of LusTRE. In particular, Figure 2 and Figure 3 compare the cross-walking measures with the cardinality-based measures. The measures are normalised between 0 and 1 to compare their behaviour. Figures 4 and 5 present the values of the cross-walking measures applied to LusTRE without any normalisation.
[Property of Quality Measures]. The analysis of results in Figure 2 and Figure 3 highlights that all the measures have different behaviours. Figure 2 shows linksets having high linkset cardinality and coverage but low average linkset importing and linkset lexical enrichment (e.g., GEMETTOEARTh), and linksets having low cardinality/coverage but high average linkset importing (e.g., as for ThISTTOGEMET). Similar observations can be taken out considering the browsing measures. Figure 3 shows linksets having high linkset cardinality and coverage but low lexical and browsing enrichment (e.g., GEMETTOEARTh), and linksets exhibiting high lexical enrichment/browsing enrichment but low linkset coverage (e.g., GEMETTOEUROVOC). Thus, to a high/low value for the cardinality-based measures does not necessarily correspond a high/low value for cross-walking measures. Moreover, cardinality-based measures do not distinguish between multilingual and browsing space improvement. In fact, to the same linkset coverage or cardinality value might correspond distinct cross-walking measure values for preferred label, alternative labels, and concepts (e.g., AGROVOCTOEUROVOC in Figure 2). The previous discrepancies suggest *the inadequacy of cardinality-based measures to evaluate the quantity and the impact of new information accessible by navigating the linkset and their difficulties in distinguishing among distinct kinds of information*.

Further considerations arise from the cross-walking measures definition and the results shown in Figures 4 and 5.

- *Average linkset measures are upper bounds for enrichment measures*. As explained in 3.2, the enrichment measures can be obtained by multiplying the linkset coverage by the average linkset measures. The linkset coverage ranges between 0 and 1, as no more than one `skos:exactMatch` link for each subject concept is allowed, while the average linkset measure is positive, thus, the average linkset measures are upper bounds for the enrichment measures. Indeed every histogram in Figure 4 is higher than the correspondent histogram in Figure 5.
- *Average linkset measures and enrichment linkset measures assess distinct qualities.* Even if each average linkset measure can be turned into its enrichment counterpart by multiplying for coverage, similar values for the average linkset qualities might correspond to quite distinct values for the enrichment measures. For example, considering AGROVOCTOEUROVOC and GEMETTOEUROVOC in Figure 4, the values of the `skos:prefLabel`-average importing are similar, about 25, whilst in Figure 5, their enrichment counterparts are very different: about 1.5 and about 9 respectively. In fact, depending on the number of subject concepts involved in the linkset, a similar amount of imported lexical labels or reached concepts per link can impact very differently on the subject thesaurus.

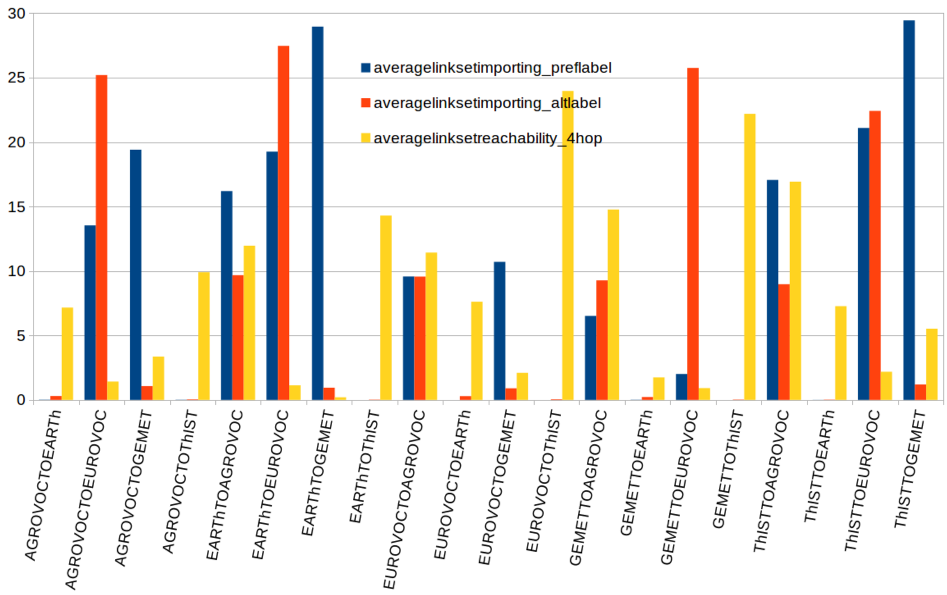

Fig. 4. Results of the average linkset reachability and the average linkset importing for `skos:prefLabel` and `skos:altLabel` applied to LusTRE.

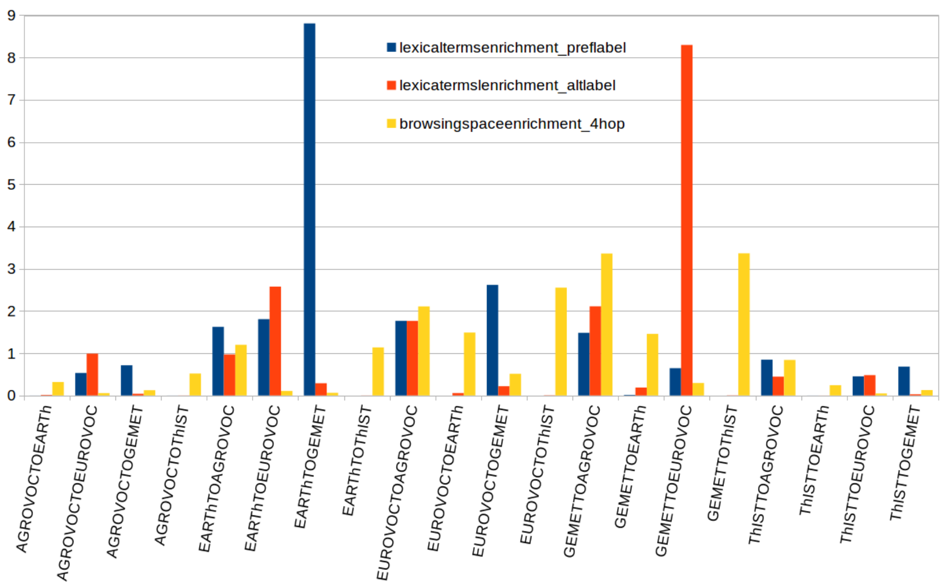

Fig. 5. Results of the browsing space and the lexical enrichment for `skos:prefLabel` and `skos:altLabel` applied to LusTRE.

[Quality of LusTRE's linksets]. Multilingualism improvement and widening of the browsing space. In Figures 4 and 5, at least one of the cross-walking measures is greater than zero for each linkset. In particular, Figure 4 shows that *every linkset brings some value, either in term of new lexical labels or concepts*, and Figure 5 shows that *these new information impact on the subject thesauri* of the linksets. LusTRE is then the classical situation in which "the union is strength": the joint use of linked thesauri provides a richer multilingualism and browsing space than using them alone.

Multilingualism improvement. The analysis identifies the linksets which bring a significant average number of newly translated labels per link and those which are the most impacting on the multilingualism of the subject thesaurus.

- *Linksets which reach a regardable number of lexical labels per link.* Figure 4 shows an exceptional average linkset importing of preferred labels for the linksets ThISTTOGEMET and EARThTOGEMET, which respectively bring 29 and 28 new `skos:prefLabel` for each link. These values can be explained by recalling some characteristics of the linked thesauri. In particular, ThIST is a bilingual thesaurus, EARTh is a bilingual thesaurus with a partial translation in Spanish, and GEMET exhibits preferred labels in more than 30 languages, as a consequence, each link from ThIST and EARTh to GEMET can gain many translations. Similar observations can be drawn for ThISTTOEUROVOC and EARThTOEUROVOC. The average linkset importing of `skos:prefLabel` for AGROVOCTOGEMET is also regardable: AGROVOC exhibits preferred labels in 23 languages, which only partially overlaps the languages exhibited by GEMET. Thus the linkset AGROVOCTOGEMET reaches one translation for each of the not overlapping languages.
  Concerning `skos:altLabel`, the topmost average linkset importings are for the linksets EARThTOEUROVOC, AGROVOCTOEUROVOC and GEMETTOEUROVOC. The fact that EUROVOC is the object thesaurus of all these topmost linksets suggests that EUROVOC is particularly rich of alternative labels. Figure 4 shows other six linksets (i.e., AGROVOCTOEUROVOC, EARThTOAGROVOC, EUROVOCTOAGROVOC, EUROVOCTOGEMET, GEMETTOAGROVOC, ThISTTOAGROVOC) that import a regardable number of preferred or alternative labels (≥ 8).
- *Linksets which improve the multilingualism of the subject thesauri.* A good average linkset importing implies that a relevant number of new lexical values are reached per link, but it does not necessarily entail an impact on the subject thesaurus multilingualism. For instance, in Figure 4, EARThTOGEMET and ThISTTOGEMET have good linkset importing values, but their lexical enrichments differ in Figure 5: it is very high for EARThTOGEMET (9 new `skos:prefLabel` per concept) while it is lower for ThISTTOGEMET (less than 1 new preferred label per concept). Nonetheless, a lexical term enrichment equal to 1 implies that the linkset has reached as many translations as the number of concepts in the subject thesaurus. In reason of that, this paper considers as impacting any lexi-

cal term enrichment equal or higher than 1. EARThTOGEMET and EUROVOCTOGEMET are the two topmost impacting linksets for `skos:prefLabel` with values close to 9 and 2.5, respectively. GEMETTOEUROVOC is the topmost impacting linkset for `skos:altLabel` with a lexical term enrichment value higher than 8. EARThTOAGROVOC, EARThTOEUROVOC, EUROVOCTOAGROVOC, GEMETTOAGROVOC, GEMETTOAGROVOC impact on both `skos:prefLabel` and `skos:altLabel` of their subject thesauri exhibiting lexical term enrichment values equal or higher than 1.

- *Linksets which poorly improve lexical labels.* Figure 4 shows a poor improvement brought by the linksets AGROVOCTOEARTh, AGROVOCTOThIST, EARThTOThIST, EUROVOCTOEARTh, EUROVOCTOThIST, GEMETTOEARTh, GEMETTOThIST, ThISTTOEARTh. These linksets poorly behave also in Figure 5 for the lexical enrichment measure. The object thesauri are EARTh and ThIST which provide a quite limited number of languages. Nonetheless, all the thesauri except ThIST improve their multilingualism either in terms of `skos:preflabel` or `skos:altlabel`, as they are subject datasets in at least one linkset with lexical term enrichment close or greater to 1.

Thesaurus browsing space widening. The following analysis identifies the linksets which bring a significant average number of new concepts per links and the most impacting on the subject browsing space.

- *Linksets which reach a significant number of concepts.* Figure 4 shows that the linksets EUROVOCTOThIST and GEMETTOThIST reach in average more than 20 new concepts per link, which is a very significant value. Other linksets with a notable average reachability, between 10 and 20, are AGROVOCTOThIST, EARThTOAGROVOC, EARThTOThIST, ThISTTOAGROVOC. The linksets AGROVOCTOEARTh, EUROVOCTOEARTh, ThISTTOEARTh, ThISTTOGEMET reach in average between 5 and 7 new concepts per link.
- *Linksets which significantly improve the browsing space of their subject thesauri*. Figure 5 shows that EARTh, EUROVOC and GEMET significantly increase their browsing spaces. In particular, considering the maximum browsing space enrichment for each subject thesaurus, EARTh duplicates the number of its concepts by traversing the linkset EARThTOAGROVOC whose browsing space enrichment is equal to 1; EUROVOC triplicates the number of its concepts by traversing the linksets EUROVOCTOAGROVOC or EUROVOCTOThIST whose browsing space enrichment is higher than 2; GEMET quadruplicates the number of its concepts by traversing the linksets GEMETTOAGROVOC or GEMETTOThIST whose browsing space enrichment is higher than 3.
- *Linksets which poorly improve the browsing space of their subject thesauri*. Considering Figures 5, the AGROVOC and ThIST exhibit a limited browsing space improvement: it is minor than 0.5 for each of the linksets where they are subject thesauri. Such as a little impact on AGROVOC and ThIST is due to a very limited

linkset coverage; Figure 3 shows that linksets having AGROVOC and ThIST as linkset subject have the lowest linkset coverage.

Summarising the discussion, cross-walking measures can evaluate information not yet considered by cardinality-based measures and they might serve to deepen the linkset quality analysis. Cross-walking measures allow evaluating the advantage in term of multilingualism and browsing space, distinguishing the number of information reached via linkset from the overall impact on the subject thesaurus.

The analysis shows that a regardable amount of lexical labels and concepts are reached via the LusTRE linksets. Thus, services and data consumers can usefully exploit the cross-walking features provided by LusTRE.

6. Conclusions and Future Work

The paper presents an approach for evaluating the quality of `skos:exactMatch` linksets among SKOS thesauri. It defines two new measures: the lexical term enrichment and the browsing space enrichment. These measures together with our previously defined average linkset importing and average linkset reachability evaluate the kind and the quantity of information accessed cross-walking the linkset. The paper compares these measures with the cardinality-based measures (linkset coverage and cardinality). All the considered measures are applied to the `skos:exactMatch` linksets available in LusTRE. The results show the measures complementarity: each measure evaluates a particular aspect of the linkset quality. The cross-walking measures can evaluate the improvements brought by the `skos:exactMatch` linksets on multilingualism and concept browsing space. The results show that the efforts spent in the creation of LusTRE linksets are paid off: the joint use of linked thesauri provide a richer multilingualism and browsing space then using the thesauri alone. Future works will extend the measures considering other kinds of linksets (e.g., `skos:relatedMatch` and `skos:broaderMatch` linksets); they will encode the quality results according to the DQV, and they will apply the proposed measures "into the wild" considering the whole set of SKOS linksets exposed in the LOD Cloud.

Acknowledgements

This research activity has been carried out within the EU project eENVplus (CIP-ICT-PSP grant No. 325232).

Notes

1. *https://tools.ietf.org/html/rfc3987*
2. *https://www.w3.org/DesignIssues/LinkedData.html*
3. *http://lod-cloud.net/*
4. *http://linkeddata.ge.imati.cnr.it/*
5. *http://www.eenvplus.eu*
6. *http://www.isprambiente.gov.it/it/pubblicazioni/pubblicazioni-di-pregio/thesaurus-italiano-di-scienze-della-terra*
7. *https://www.eionet.europa.eu/gemet/*

8. *http://eurovoc.europa.eu/*
9. *http://lod-cloud.net/*
10. *https://www.w3.org/2001/sw/wiki/Skosify*
11. *https://www.poolparty.biz/*
12. *http://silkframework.org/*
13. *http://aksw.org/Projects/LIMES.html*
14. *http://showcase.eenvplus.eu/client/editor.htm*
15. *http://plugins.qgis.org/plugins/qsphere/*
16. *http://showcase.eenvplus.eu/client/geoportal.htm*

References

Acosta, M., Zaveri, A., Simperl, E., Kontokostas, D., Auer, S. and Lehmann, J. (2013), Crowdsourcing linked data quality assessment, *in* 'The Semantic Web - ISWC 2013, Sydney, NSW, Australia, October 21-25, 2013, Proceedings, Part II', pp. 260–276. **URL:** *https://doi.org/10.1007/978-3-642-41338-4_17*

Al-Hakim, L. (2006), *Challenges of Managing Information Quality in Service Organizations*, Gale virtual reference library, Idea Group Pub.

Albertoni, R., De Martino, M., Di Franco, S., De Santis, V. and Plini, P. (2014*b*), 'EARTh: An environmental application reference thesaurus in the linked open data cloud', *Semantic Web* **5**(2), 165–171. **URL:** *https://doi.org/10.3233/SW-130122*

Albertoni, R., De Martino, M. and Podestà, P. (2014*a*), Environmental thesauri under the lens of reusability, *in* 'EGOVIS 2014', Vol. 8465 of *Lecture Notes in Computer Science*, Springer, pp. 222–236. **URL:** *https://doi.org/10.1007/978-3-319-10178-1_18*

Albertoni, R., De Martino, M. and Podestà, P. (2015), A linkset quality metric measuring multilingual gain in SKOS thesauri, *in* A. Rula, A. Zaveri, M. Knuth and D. Kontokostas, eds, 'Proc. of the 2nd Workshop on Linked Data Quality co-located with ESWC 2015, Portorož, Slovenia, June 1, 2015.', Vol. 1376 of *CEUR Workshop Proceedings*, CEUR-WS.org. **URL:** *http://ceur-ws.org/Vol-1376*

Albertoni, R., De Martino, M. and Podestà, P. (2016), Linkset quality assessment for the thesaurus framework LusTRE, *in* Garoufallou et al. (2016), pp. 27–39. **URL:** *https://doi.org/10.1007/978-3-319-49157-8*

Albertoni, R., De Martino, M., Podestà, P., Abecker, A., Wössner, R. and Schnitter, K. (2018), 'LusTRE: a framework of linked environmental thesauri for metadata management', *Earth Science Informatics*. To appear, **URL:** *https://doi.org/10.1007/s12145-018-0344-8*

Albertoni, R. and Gómez Pérez, A. (2013), Assessing linkset quality for complementing third-party datasets, *in* Proceedings of the Joint EDBT/ICDT 2013 Workshops, EDBT '13, ACM, New York, NY, USA, pp. 52–59. *http://doi.acm.org/10.1145/2457317.2457327*

Albertoni, R. and Isaac, A. (2016), Data on the web best practices: Data quality vocabulary, W3C note, W3C. **URL:** *https://www.w3.org/TR/vocab-dqv/.*

Alexander, K., Zhao, J., Cyganiak, R. and Hausenblas, M. (2011), Describing linked datasets with the VoID vocabulary, W3C note, W3C. **URL:** *http://www.w3.org/TR/void/*

Baker, T., Caracciolo, C., Doroszenko, A. and Suominen, O. (2016), GACS core: Creation of a global agricultural concept scheme, *in* Garoufallou et al. (2016), pp. 311–316. **URL:** *https://doi.org/10.1007/978-3-319-49157-8_27*

Bizer, C. and Cyganiak, R. (2009), 'Quality-driven information filtering using the WIQA policy framework', *J. Web Sem.* **7**(1), 1–10. **URL:** *http://doi.org/10.1016/j.websem.2008.02.005*

Bizer, C., Heath, T. and Berners-Lee, T. (2009), 'Linked data - the story so far', *Int. J. Semantic Web Inf. Syst.* **5**(3), 1–22. **URL:** *http://doi.org/10.4018/jswis.2009081901*

Calegari, N., Loscio, B. F. and Burle, C. (2017), Data on the web best practices, W3C recommendation, W3C. **URL:** *https://www.w3.org/TR/dwbp/*

Caracciolo, C., Stellato, A., Morshed, A., Johannsen, G., Rajbhandari, S., Jaques, Y. and Keizer, J. (2013), 'The AGROVOC linked dataset', *Semantic Web* **4**(3), 341–348. **URL:** *https://doi.org/10.3233/SW-130106*

De Martino, M. and Albertoni, R. (2011), 'A multilingual/multicultural semantic-based approach to improve data sharing in a SDI for nature conservation', *IJSDIR* **6**, 206–233. **URL:** *http://ijsdir.jrc.ec.europa.eu/index.php/ijsdir/article/view/202*

Debattista, J., Auer, S. and Lange, C. (2016*a*), Luzzu - A framework for linked data quality assessment, *in* 'Tenth IEEE International Conference on Semantic Computing, ICSC 2016, Laguna Hills, CA, USA, February 4-6, 2016', IEEE Computer Society, pp. 124–131. **URL:** *https://doi.org/10.1109/ICSC.2016.48*

Debattista, J., Auer, S. and Lange, C. (2016*b*), 'Luzzu - A methodology and framework for linked data quality assessment', *J. Data and Information Quality* **8**(1), 4:1–4:32. **URL:** *https://doi.org/10.1145/2992786*

Färber, M., Bartscherer, F., Menne, C. and Rettinger, A. (2018), 'Linked data quality of dbpedia, freebase, opencyc, wikidata, and YAGO', *Semantic Web* **9**(1), 77–129. **URL:** *https://doi.org/10.3233/SW-170275*

Garoufallou, E., Coll, I. S., Stellato, A. and Greenberg, J., eds (2016), *Metadata and Semantics Research - 10th International Conference, MTSR 2016, Göttingen, Germany, November 22-25, 2016, Proceedings*, Vol. 672 of *Communications in Computer and Information Science*. **URL:** *https://doi.org/10.1007/978-3-319-49157-8*

Guéret, C., Groth, P. T., Stadler, C. and Lehmann, J. (2012), Assessing linked data mappings using network measures, *in* E. Simperl, P. Cimiano, A. Polleres, Ó. Corcho and V. Presutti, eds, 'ESWC', Vol. 7295 of *Lecture Notes in Computer Science*, Springer, pp. 87–102. **URL:** *http://doi.org/10.1007/978-3-642-30284-8_13*

Harris, S. and Seaborne, A. (2013), SPARQL 1.1 query language, W3C recommendation, W3C. **URL:** *http://www.w3.org/TR/sparql11-query/*

Kontokostas, D., Westphal, P., Auer, S., Hellmann, S., Lehmann, J., Cornelissen, R. and Zaveri, A. (2014), Test-driven evaluation of linked data quality, *in* 'Proceedings of the 23rd International Conference on World Wide Web', WWW '14, ACM, New York, NY, USA, pp. 747–758. **URL:** *https://doi.org/10.1145/2566486.2568002*

Lacasta, J., Falquet, G., Zarazaga-Soria, F. J. and Nogueras-Iso, J. (2016), 'An automatic method for reporting the quality of thesauri', *Data Knowl. Eng.* **104**, 1–14. **URL:** *https://doi.org/10.1016/j.datak.2016.05.002*

Mader, C., Haslhofer, B. and Isaac, A. (2012), Finding quality issues in skos vocabularies, *in* P. Zaphiris, G. Buchanan, E. Rasmussen and F. Loizides, eds, 'TPDL', Vol. 7489 of *Lecture Notes in Computer Science*, Springer, pp. 222–233. **URL:** *http://doi.org/10.1007/978-3-642-33290-6_25*

Miles, A. and Bechhofer, S.(2009), SKOS simple knowledge organization system reference, W3C recommendation, W3C. **URL:** *http://www.w3.org/TR/skos-reference*

Pipino, L., Lee, Y. W. and Wang, R. Y. (2002), 'Data quality assessment', *Commun. ACM* **45**(4), 211–218. **URL:** *http://doi.org/10.1145/505248.506010*

Quarati, A., Albertoni, R. and De Martino, M. (2017), 'Overall quality assessment of SKOS thesauri: An AHP-based approach', *J. Information Science* **43**(6), 816–834. **URL:** *https://doi.org/10.1177/0165551516671079*

Radulovic, F., Mihindukulasooriya, N., García-Castro, R. and Gómez-Pérez, A. (2018), 'A comprehensive quality model for linked data', *Semantic Web* **9**(1), 3–24. **URL:** *https://doi.org/10.3233/SW-170267*

Schreiber, G. and Raimond, Y. (2014), RDF 1.1 primer, W3C note, W3C. **URL:** *https://www.w3.org/TR/rdf11-primer/*

Suominen, O. and Hyvönen, E. (2012), Improving the quality of skos vocabularies with skosify, *in* A. ten Teije, J. Völker, S. Handschuh, H. Stuckenschmidt, M. d'Aquin, A. Nikolov, N. Aussenac-Gilles and N. Hernandez, eds, 'EKAW', Vol. 7603 of *Lecture Notes in Computer Science*, Springer, pp. 383–397. **URL:** *http://doi.org/10.1007/978-3-642-33876-2_34*

Suominen, O. and Mader, C. (2014), 'Assessing and improving the quality of skos vocabularies', *J. Data Semantics* **3**(1), 47–73. **URL:** *http://doi.org/10.1007/s13740-013-0026-0*

Wang, R. Y. and Strong, D. M. (1996), 'Beyond accuracy: What data quality means to data consumers', *J. Manage. Inf. Syst.* **12**(4), 5–33. **URL:** *https://doi.org/10.1080/07421222.1996.11518099*

Zaveri, A., Kontokostas, D., Sherif, M. A., Bühmann, L., Morsey, M., Auer, S. and Lehmann, J. (2013), User-driven quality evaluation of dbpedia, *in* 'I-SEMANTICS 2013 Graz, Austria, September 4-6, 2013', pp. 97–104. **URL:** *https://doi.org/10.1145/2506182.2506195*

Zaveri, A., Rula, A., Maurino, A., Pietrobon, R., Lehmann, J. and Auer, S. (2016), 'Quality assessment for linked data: A survey', *Semantic Web* **7**(1), 63–93. **URL:** *https://doi.org/10.3233/SW-150175*