# Enhancing the Discoverability and Interoperability of Multi-disciplinary Semantic Repositories

Doron Goldfarb[1[0000-0003-1183-6041]] and Yann Le Franc[2[0000-0003-4631-418X]]

[1] Environment Agency Austria, Vienna, Austria
[2] e-Science Data Factory, Paris, France
doron.goldfarb@umweltbundesamt.at, ylefranc@esciencefactory.com

**Abstract.** The aggregation of multi-disciplinary information is a challenge faced by large-scale data infrastructures serving scientific domains such as biodiversity, agronomy or ecology. This requires the integration of ontologies or thesauri from different domains. These semantic resources are often hosted within domain specific repositories which can be harvested for that purpose. The lack of discoverability, the technical and metadata heterogeneity of the semantic repositories pose a challenge for their effective integration. In this context, we argue that there is a need for a semantic lookup-service to access and use this heterogeneous landscape. We then present a proof-of-concept design and implementation for harvesting different ontology repositories (BioPortal, AgroPortal and EBI-OLS). We show some preliminary analytics and discuss technical issues regarding aggregation. Finally, we conclude with an open call for collaboration to address the issues hampering such initiatives.

**Keywords:** Ontology libraries, Semantic annotation, Ontology lookup service, EUDAT.

## 1 Introduction

Semantic technologies are increasingly used by domain-specific Research Infrastructures (RIs) and large-scale multi-disciplinary infrastructures such as EUDAT[1]. Semantically-enabled services offer a framework to aggregate data from multiple sources, enhancing discoverability and interoperability. The EUDAT pilot service B2Note[2] is one such service, allowing the creation of semantic annotations of datasets within and outside of the EUDAT infrastructure. The process of annotation is about "attach[ing] data to some other piece of data" [1]. In the scope of the Semantic Web, this usually refers to the contextualisation of information within a wider knowledge graph in order to support discovery and, eventually, automated reasoning. Such a vision can only be made possible through the wide-spread and repeated annotation with concepts defined in ontologies, thesauri or taxonomies. Throughout this work,

---

[1] http://www.eudat.eu/
[2] https://eudat.eu/news/annotate-your-research-data-with-b2note

we refer to such formalised knowledge representation structures as semantic resources, without any consideration for their format.

Providing domain specific concepts within a multi-disciplinary infrastructure requires their discovery and aggregation from different semantic resources available throughout the Web. This is also particularly true for RIs in the domain of biodiversity and ecology, where biology is linked to heterogeneous fields such as chemistry, molecular biology and earth science. In recent years, however, the number of available semantic resources has steadily grown to an extent making it hard to maintain the overview on "what's out there" and to identify the locations where they can be retrieved.

Dedicated repositories have thus been conceived to extend the discoverability of semantic resources by providing single access points for retrieving information about and from multiple, usually domain specific, semantic resources. Called "ontology libraries" by d'Aquin and Noy [2], these semantic repositories often provide programmatic access as (REST) API or via query languages such as SPARQL. They can thus be used to identify available semantic resources and moreover usually offer the advantage of hosting them in a homogenised form. This includes structured descriptive metadata about a resource such as name, acronym and version, as well as homogeneous extracts of its content which usually encompasses information about concepts and related properties.

Harvesting content from different domain specific semantic repositories can therefore support the aggregation of domain specific concepts for the semantic annotation of multi-disciplinary content. This endeavour, however, still remains a challenge as the large number of available semantic repositories raises the problem of their discoverability and interoperability.

In the context of EUDAT we designed a proof-of-concept service to aggregate multi-disciplinary semantic resources. This Semantic Lookup Service shall periodically retrieve the content from a set of registered semantic repositories and feed the results into a search index supporting concept discovery and auto-completion, used by the data annotation service B2Note. The development of such a centralised platform will increase the discoverability of the existing resources for the domain knowledge experts and for the growing eco-system of semantic tools, supporting the re-use of the semantic resources. Furthermore, the aggregation of content from large numbers of semantic repositories enables various types of analysis and metrics.

We argue that such a service will be of benefit especially to the life-sciences domain, since a huge proportion of existing repositories, such as BioPortal [3] and EBI-OLS [4], is rooted there, reflecting the already established tradition of using semantic resources. Providing a consolidated view on the semantic resource landscape present there will support related researchers but also foster the re-use of their resources in related domains such as agronomy, biodiversity and ecosystems research. In the context of the latter for example, initiatives such as the ILTER (International Long Term Ecosystem Research) network increasingly employ semantic resources such as the Environmental Thesaurus [5] for facilitating the harmonisation of heterogeneous data from its members [6]. The establishment of a repository dedicated to such resources is planned [7] and will augment already existing initiatives such as AgroPortal [8] for the Agronomy/Agrology domain. The aggregation of these repositories provides the opportunity to identify cross-domain overlaps in terminology, potentially leading to mutual re-use and better cross-domain interoperability.

The remainder of this paper is structured as follows. Section 2 provides an overview on related work, identifies the key challenges and argues for the need for harmonisation between the existing solutions. It is then followed by the description of the design and first implementation of the proof-of-concept Semantic Lookup Service and our initial approach to harvesting concepts hosted in different semantic repositories in Section 3. Section 4 features a discussion of the results while section 5 gives an outlook for future work.

## 2 Related Work and Challenges

Semantic repositories seek to offer unique software platforms for extending the discoverability of semantic resources. Several implementations and approaches to semantic repositories have been developed and a first classification was proposed in 2012 by d'Aquin et al. [2]. While some of the approaches mentioned by the authors appear to have stalled, many others have emerged such as for example the SKOS oriented FINTO[3] service which is based on the SKOSMOS framework [9], the generic and curated Linked Open Vocabulary platform [10], the ANDS Research Vocabulary service[4] based on SISSVOC [11], and Ontobee [12]. In the biomedical domain, projects such as the above mentioned BioPortal and EBI-OLS evolved into advanced repositories. BioPortal has been reused in a growing variety of domains including Agriculture (AgroPortal) and Earth Sciences[5]. All these platforms offer means to harvest their content via RESTful APIs or SPARQL endpoints and thus support the aggregation of their content. The current variety of technical solutions increases the choice of offered functionalities but comes with burdens for interoperability. Indeed, the comparison of the different repositories revealed a large metadata and API heterogeneity. This represents a challenge to aggregate these resources into a multidisciplinary semantic index.

Another major challenge is the discoverability of semantic resources. Indeed, the increasing number of semantic repositories makes it difficult to find all of them. In addition, many resources are not necessarily registered in a repository and can be rarely found via Google searches, leading to the situation that they are only known and (re-)used within a specific community.

We identified three main needs to address these challenges: a common metadata description, a common framework for API interoperability and a central hub to access the wealth of semantic resources.

The need for a common metadata description for semantic resources and semantic repositories has been identified by several initiatives such as the OBO Foundry [13]

[3] http://finto.fi/en/
[4] https://vocabs.ands.org.au/
[5] http://semanticportal.esipfed.org/

for the bio-medical domain, the ontology metadata schema proposed by LOV, and the Ontology Metadata Vocabulary [14]. One of the key challenges is to find a consensus between these different initiatives and to define a unique minimal common metadata set in order to enhance the interoperability between the different existing resources.

The problem of API interoperability is a generic problem for interlinking web services infrastructure. In the past few years, different initiatives have emerged to address this issue, including, for web-based APIs, the W3C HYDRA working community [15], the OpenAPI Initiatives[6] and the smart API [16], or, for RDF based datasets and resources, the W3C Vocabulary of Interlinked Datasets[7] and Linked Data Fragments[8]. The design of a central hub for API-providing repositories represents a unique opportunity to test and benchmark the different approaches to API interoperability.

The vision of a centralised service for discovering, searching, exploring and reusing semantic resources and related documents has already been proposed by several initiatives. Semantic search engines such as Swoogle [17], FalconS [18] or Watson [19] aimed at crawling and mining the web for semantic resources and offered means to search the results. Although they became valuable resources for knowledge workers, these different initiatives appear to have been discontinued. Other approaches sought to provide distributed search facilities across semantic repositories, such as the "Network of Ontology Repositories" [20], OntoCAT [21] and OntoHub [22].

In contrast to the latter, the approach presented in this paper aims at proposing a centrally aggregated search index which is not limited to locally stored resources but includes concept level extracts from remotely harvested semantic repositories. This index can then be used as a semantic search engine based on - in contrast to Swoogle and related work - registered resources and repositories instead of Web crawling.

We believe that such an approach is beneficial for the quality of the content and its re-use. Moreover, it would provide means for large-scale analysis of the different resources such as the recent analysis performed in BioPortal [23] and provide meaningful information to ontologists, data scientists and knowledge engineers.

## 3 The Semantic Look Up Service

Based on the need identified in the previous sections, we designed an initial proof-of-concept service for cataloguing semantic resources. In this section we first describe the general design principles, then the current implementation.

The design of the service followed several identified requirements, centred on providing means to adequately describe and register semantic repositories in a way that the relevant information about the hosted resources and the concepts therein could be mapped to a common representation for indexing. Stored in a database, such descriptions should enable a harvesting service to retrieve the content at regular intervals, transforming and storing it as common index representation. This basic infrastructure should be flexible enough to allow the provision of additional services such

[6] https://www.openapis.org/
[7] https://www.w3.org/TR/void/
[8] http://linkeddatafragments.org/

as a public catalogue of described repositories/resources and an API for data analysis. A schematic of the proposed architecture and the data flow between the different elements of the service is shown in Figure 1.

The service should be composed of 5 main components: (1) a web interface to capture the description of external semantic repositories provided by their managers, (2) a database storing the repository descriptions used by (3) an information harvester collecting the descriptions to harvest the repository contents via their API and store them into the database. This information will be used to build and update (4) an index of all the concepts contained in the different registered repository for fast retrieval and use for semantic services. Finally, another web interface (5) should be designed to discover, visualise and interact with the catalogue, providing views of registered semantic repositories and the ontologies and vocabularies harvested, including information that can be gathered and extracted from them, such as inter-repository overlap both on resource and on concept level.

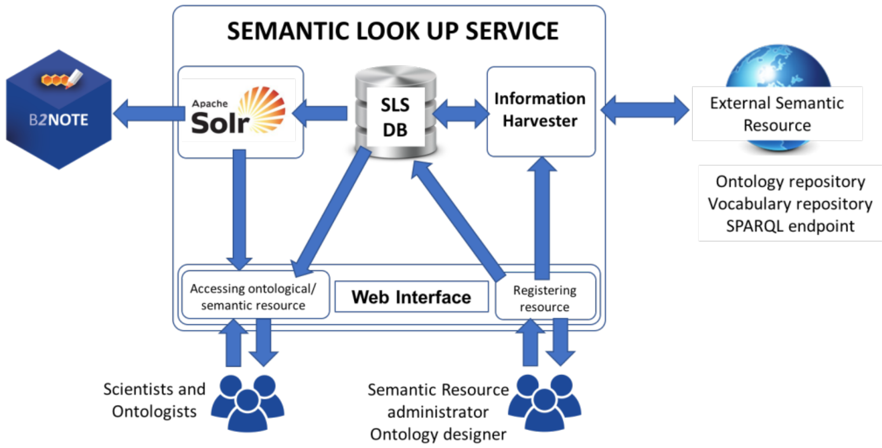


**Fig. 1.** Proposed architecture and data flow in EUDAT's Semantic Lookup Service

### 3.1 Current Implementation

We developed an initial implementation based on Python and command line scripts to provide an initial index of bio-medical semantic concepts for B2NOTE. We created custom scripts to harvest 5 million concepts from 494 ontologies hosted in BioPortal which populated a first instance of a SolR[9] index. The initial SolR schema has a focus on information about terms/classes defined/reused in the individual resources, but our approach can easily be extended to also cover other aspects such as properties. Table 1 lists the respective index fields. Besides fields describing concepts on individual and resource level, additional ones are conceived for filtering/organising search results. One such field is dedicated to listing concept reuse across resources and can support

[9] http://lucene.apache.org/solr/

the ranking of results, while another one provides information about the domain(s) the concept belongs to which can be used for limiting the search space. While the former can be automatically derived by analysing the harvested resources for conceptual overlaps, the latter should be provided by the repositories themselves and is currently only rarely available.

**Table 1.** Fields used for the SOLR lookup index

| | |
|---|---|
| Concept IRI | IRI of the concept. |
| Concept Label | Human readable label of the concept. |
| Concept Description | Definition of the concept. |
| Concept Short_form | Short form of the concept. |
| Concept Synonyms | List of synonym labels referenced for the concept. |
| Resource Acronym | Acronym of the resource the concept pertains to. |
| Resource IRI | IRI of the resource the concept pertains to. |
| Resource name | Name of the resource the concept pertains to. |
| Resource vdate | Resource "released" field information. |
| Resource version | Resource "version" field information. |
| Acrs_of_resources_reusing_uri | List of acronyms for the resources reusing the concept. |
| Domains (not harvested yet) | Scientific domain covered by the resource |

The initial workflow for harvesting the BioPortal API was directly coded as a Python script. This "plug-in" based approach is clearly not scalable and requires building scripts for every repository and maintaining them accordingly. We thus concentrated on developing a more efficient and generic approach for harvesting different repositories, focusing our effort on harvesting REST APIs and leaving the harvesting of SPARQL endpoints for future work.

To acquire the information needed for the SolR index (see Table 1), none of the analysed platforms, i.e. BioPortal, EBI-OLS and AgroPortal, provide one single function for retrieving the full set of fields and their harvesting thus involves several steps. We identified a two-step pattern to access this information. For each repository, an initial request retrieves basic resource level information and for each of the retrieved resources, additional requests then acquire information both on resource as well as on concept level.

Our initial approach to provide a common description framework for the three observed repositories uses a JSON description of the query sequence identified above. It contains information about the necessary query URLs as well as the locations of the data in the response sets from the different APIs, mapped to the respective fields of the SOLR index via JSONPath[10] expressions. Since a more detailed description is beyond the scope of this paper, it will be provided in a separate work. As shown in the next section, however, this approach enables us to successfully re-use one base implementation across three repositories, one of which having quite a different API implementation compared to the other two.

## 4 Analysing three Semantic Repositories

We applied our JSON/JSONPath based harvesting description to three existing repositories, BioPortal and its derivative AgroPortal, as well as EBI-OLS, this section provides results of a preliminary analysis. We were able to retrieve 96% of the available concepts (13,660,813 out of 14,226,183) from 93% of the semantic resources (786 out of 843). A first analysis of this initial dataset showed that 8,840,852 concepts were unique when distinguished by their URI and 6,109,756 when compared by strictly matching preferred labels. Table 2 provides data for each repository. Included in parentheses is the number of encountered resources having at least one concept vs. the stated number of resources hosted by each repository. The lower numbers for resources with concepts is due to the fact that missing ones are registered as private, as summary description only, feature only properties but no concepts, or are SKOS based resources. The latter - including a large fraction of AgroPortal, e.g. AGROVOC - were not harvested as they are considered as instances and not classes in the repository. We will investigate this in more detail and seek to extend our harvester in this regard.

**Table 2.** Total and unique number of concepts (no instances) present in three repositories

| | AgroPortal (63/64) | BioPortal (534/586) | EBI-OLS (189/193) |
|---|---|---|---|
| Total | 1,198,472/1,200,845 | 7,569,311/8,130,580 | 4,893,030/4,894,758 |
| Unique URI | 1,186,681 | 6,659,704 | 4,235,425 |
| Unq. Label | 1,122,242 | 5,379,485 | 3,938,468 |

Checking for inter-repository overlap on resource level, we compiled an alignment of the resource descriptions retrieved from each API. Noting some ambiguities regarding resource acronyms and discrepancies regarding resource URIs, we aligned the resources by their name. This operation still required some manual editing to compen-

[10] http://goessner.net/articles/JsonPath/

sate for encountered naming discrepancies. The assessed resource overlap is presented in Figure 2, updating and extending the comparison between BioPortal and EBI-OLS as of 2011 [21]. A strong overlap between EBI-OLS (grown by about ⅓ since 2011) and BioPortal (Almost tripled since 2011) becomes immediately visible, their common 131 resources (113 of them OBO foundry related) now represent 67.9% of the EBI-OLS resources (45.5% in 2011). Another interesting observation is that on resource level, AgroPortal has higher overlap with EBI-OLS than with BioPortal. This is due to a set of crop specific ontologies taken from the Crop Ontology[11] project, hosted in both EBI-OLS and AgroPortal but not in BioPortal. Besides the identified overlaps, each repository provides a unique set of resources not present in the others.

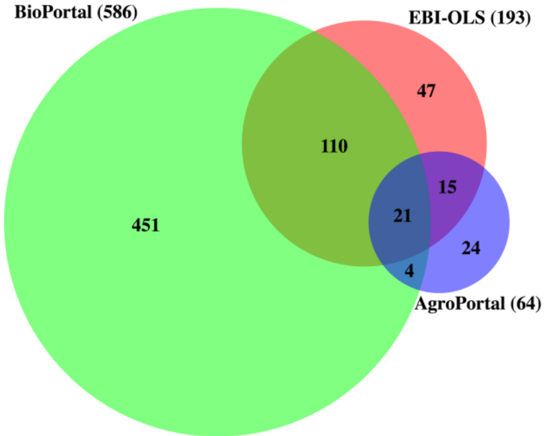


**Fig. 2.** Resources shared between EBI-OLS, BioPortal and AgroPortal

As stated above, the comparison of the resource descriptions revealed some relatively rare (affecting 36 out of 786 resources investigated) but nevertheless notable ambiguities regarding resource acronyms. We found (1) similar resources to have different acronyms in different repositories, such as the "Beta Cell Genomics Ontology" having "bcgo" in EBI-OLS and "obi_bcgo" in BioPortal, and (2) different resources to share similar acronyms across different repositories, such as "aeo" standing for "Anatomical Entity Ontology" in EBI-OLS/BioPortal and for "Agricultural Experiments Ontology" in AgroPortal.

Our observations show that acronyms are not always uniquely assigned to the same resources across repositories and such discrepancies can also be found in resource names. These ambiguities in our opinion provide a rather strong incentive to seek for a global name authority for semantic resources. We note, however, that establishing unique resource prefixes across different domains and communities could involve significant effort such as changing existing identifiers, which, as stated in [24], might outweigh the intended benefit. Given that we encountered relatively few such cases in

[11] http://www.cropontology.org

the observed repositories nevertheless suggests that the consultation of services such as prefixcommons[12] should be integrated in existing repositories.

For the further analysis of the aggregated resources, future work will concentrate on mappings at concept level, considering algorithms such as LOOM [25] and other approaches already employed in BioPortal[13] and other initiatives.

## 5 Conclusions & Outlook

In this paper we argue that there is a clear need for a centralised semantic look up service allowing to aggregate multi-disciplinary semantic resources. We emphasised that this effort is hampered by the lack of a common metadata set to describe the semantic resources, API interoperability and discoverability of the existing resources. We described our initial approach to build an index of multi-disciplinary concepts for semantically-enabled services in EUDAT. Our work presents the design of an initial proof-of-concept enabling different stakeholders to start referencing the different resources and serving as testbed for different solutions to aggregating multiple semantic repositories. We are presenting here our current implementation and the initial harvesting experiments that were performed. These experiments show that the centralised aggregation of multiple repositories also enables cross-repository analysis which is useful for studying the present landscape and improving data quality and thus interoperability.

In the future, we will extend the number of repositories to propose a general description of the harvesting workflows and we will align this work with the existing state-of-the art approaches for API interoperability. In parallel we will design and build an initial web interface to further extend the number of repositories and capture mapping information between their internal data model and the information needed for the SolR index. Finally, we will work on improving the SolR index by adding filters and facets to provide more usable search and exploration facilities across millions of terms.

We strongly believe that this effort can only be achieved through an extensive international collaboration between the different semantic repositories, the different initiatives proposing metadata representation of the semantic resources and the initiative working on API interoperability. Such collaboration has been discussed and initiated during different events organised in the context of EUDAT[14] and in collaboration with LifeWatch Italy[15]. Since these topics are in line with the general scope of the RDA Vocabulary and Semantic Service Interest Group[16] we are now working in the

---

[12] https://prefixcommons.org

[13] https://www.bioontology.org/wiki/index.php/BioPortal_Mappings

[14] https://www.eudat.eu/events/trainings/co-located-eudat-semantic-working-group-workshop-9th-rda-plenary-barcelona-3-4

[15] http://www.servicecentrelifewatch.eu/ontology-semantic-web-for-biodiversity-ecosystem-research

[16] https://www.rd-alliance.org/groups/vocabulary-services-interest-group.html

context of this interest group and hope to raise the interest on a global level and work in alignment with similar initiatives such as OntoHub and OntoCAT.

**Acknowledgements**

This work has been supported by EUDAT, funded by the European Union under the Horizon 2020 programme - DG CONNECT e-Infrastructures (Contract No. 654065). The authors would like to thank the participants of the different organised workshop for feedback on this project and their active contribution. We would also like to thank B. Magagna for her invaluable contribution, S. Cox and the RDA Vocabulary Service Interest Group for their interest and support.

## References

1. Oren, E., Möller, K., Scerri, S., et al.: What are semantic annotations?. DERI, Galway (2006).
2. d'Aquin, M. and Noy, N. F.: Where to publish and find ontologies? A survey of ontology libraries. Web Semantics: Science, Services and Agents on the World Wide Web, 11, pp. 96–111 (2012).
3. Whetzel, P. L., Noy, N. F. and Shah, N. H.: BioPortal: enhanced functionality via new Web services from the National Center for Biomedical Ontology to access and use ontologies in software applications. Nucleic Acids Research, 39(Web Server issue), (2011).
4. Jupp, S., Burdett, T., Leroy, C. and Parkinson, H.: A new Ontology Lookup Service at EMBL-EBI. In: Malone, J. et al. (eds.) Proceedings of SWAT4LS International Conference 2015, pp. 118-119, CEUR-WS, vol. 1546, ceur-ws.org, online (2015).
5. Schentz, H., Peterseil, J. and Bertrand, N.: EnvThes-interlinked thesaurus for long term ecological research, monitoring, and experiments. In: Page, B. et al (eds.) Proceedings of the 27th Int'l Conference on Environmental Informatics, Shaker, Herzogenrath (2013).
6. Vanderbilt, K. L., Lin, C.-C., Lu, S.-S., et al.: Fostering ecological data sharing: collaborations in the International Long Term Ecological Research Network, Ecosphere, 6(10), pp. 1–18 (2015).
7. Fiore, N., Magagna, B. and Goldfarb, D.: EcoSemanticPortal: facilitating discovery and interoperability of ontologies and thesauri in the ecological domain. In: Algergawy, A. et al. (eds.) Proc. of the 2nd International Workshop on Semantics for Biodiversity, (In Press).
8. Jonquet, C., Toulet, A., Arnaud, E., et al.: Reusing the NCBO BioPortal technology for agronomy to build AgroPortal. In: Jaiswal, P. et al. (eds.) Proceedings of the 7th International Conference on Biomedical Ontologies (ICBO-BioCreative), CEUR-WS, vol. 1747, ceur-ws.org, online (2016).
9. Suominen, O., Ylikotila, H., Pessala, S., et al.: Publishing SKOS vocabularies with Skosmos. Article manuscript submitted for review, http://skosmos.org/publishing-skos-vocabularies-with-skosmos.pdf, last accessed 2017/09/18.
10. Vandenbussche, P.-Y., Atemezing, G. A., Poveda-Villalón, M. and Vatant, B.: Linked Open Vocabularies (LOV): a gateway to reusable semantic vocabularies on the Web. Semantic Web, 8(3), pp. 437–452 (2017).
11. Cox, S., Yu, J. and Rankine, T.: SISSVoc: A Linked Data API for access to SKOS vocabularies. Semantic Web, 7(1), pp. 9-24 (2016).

12. Xiang, Z., Mungall, C., Ruttenberg, A. and He, Y.: Ontobee: A Linked Data Server and Browser for Ontology Terms. In: Bodenreider, O. et al. (eds.) Proceedings of the 2nd International Conference on Biomedical Ontologies (ICBO), CEUR-WS, vol. 833, pp. 279-281, CEUR-WS.org, online (2011).
13. Smith, B., Ashburner, M., Rosse, et al.: The OBO Foundry: coordinated evolution of ontologies to support biomedical data integration. Nature Biotechnology, 25(11), pp. 1251 - 1255 (2007).
14. Hartmann, J., Palma, R., Sure, Y., et al.: Ontology Metadata Vocabulary and Applications, In: Meersman R. et al. (eds.) On the Move to Meaningful Internet Systems 2005: OTM 2005 Workshops. OTM 2005. LNCS, vol. 3762, pp. 906-915, Springer, Heidelberg (2005).
15. Lanthaler, M. and Gütl, C.: Hydra: A Vocabulary for Hypermedia-Driven Web APIs. In: Bizer, C. et al. (eds.) Proceedings of the WWW2013 Workshop on Linked Data on the Web, CEUR-WS, vol. 996, ceur-ws.org, online (2013)
16. Verborgh, R. and Dumontier, M.: A Web API ecosystem through feature-based reuse. arXiv preprint, arXiv:1609.07108, arXiv (2016).
17. Finin, T., Ding, L., Pan, R., et al.: Swoogle: Searching for knowledge on the semantic web. In: Cohn, A. (ed.) Proceedings of the 20th National Conference on Artificial Intelligence, pp. 1682–1683, AAAI Press, Palo Alto (2005).
18. Cheng, G., Ge, W. and Qu, Y.: Falcons: searching and browsing entities on the semantic web. In: Huai, J., et al. (eds.) Proceedings of the 17th international conference on World Wide Web, pp. 1101–1102, ACM, New York (2008).
19. d'Aquin, M., Baldassare, C., Gridinoc, L., et al.: Watson: Supporting next generation semantic web applications. In: Nunes, M. P., et al. (eds.) Proceedings of the IADIS International Conference on WWW/Internet. IADIS Press, (2007).
20. Viljanen, K., Tuominen, J. and Hyvönen, E.: A Network of Ontology Repositories, https://seco.cs.aalto.fi/publications/2010/viljanen-et-al-onki-nor-2010.pdf (2010), last accessed 2017/07/28.
21. Adamusiak, T., Burdett T., Kurbatova, N., et al.: OntoCAT -- simple ontology search and integration in Java, R and REST/JavaScript, BMC Bioinformatics, 12, 218 (2011).
22. Mossakowski, T., Kutz, O. and Codescu, M.: Ontohub - a repository engine for heterogeneous ontologies and alignments, http://ontolog.cim3.net/file/work/OpenOntologyRepository/Ontohub/ontohub--TillMossakowski-et-al_20130621a.pdf, last accessed 2017/07/28.
23. Kamdar, M. R., Tudorache, T. and Musen, M. A.: A Systematic Analysis of Term Reuse and Term Overlap across Biomedical Ontologies. Semantic Web, (preprint), IOS Press (2017).
24. McMurry, J. A., Juty, N., Blomberg, N., et al.: Identifiers for the 21st century: How to design, provision, and reuse persistent identifiers to maximize utility and impact of life science data. PLoS Biology, 15(6), (2017).
25. Ghazvinian A., Noy, N. F., and Musen, M. A.: Creating Mappings for Ontologies in Biomedi-cine: Simple Methods Work. In: AMIA Annual Symposium Proceedings, pp. 198-202, AMIA, online (2009).