VocBench 3: a Collaborative Semantic Web Editor for Ontologies, Thesauri and Lexicons

Armando Stellato[a], Manuel Fiorelli[a], Andrea Turbati[a], Tiziano Lorenzetti[a], Willem van Gemert[b], Denis Dechandon[b], Christine Laaboudi-Spoiden[b], Anikó Gerencsér[b], Anne Waniart[b], Eugeniu Costetchi[b], Johannes Keizer[c]

[a] *University of Rome, Tor Vergata, Via del Politecnico 1, 00133 Rome, Italy*

[b] *Publications Office of the European Union, Information Management Directorate, Standardisation Unit, Metadata Sector, 2985 Luxembourg, LUXEMBOURG*

[c] *GODAN secretariat, c/o CABI Head Office, Nosworthy Way, Wallingford, Oxfordshire, OX10 8DE, UK*

Abstract. VocBench is an open source web platform for the collaborative development of datasets complying with Semantic Web standards. Since its public release – five years ago – as an open source platform, VocBench has attracted a growing user community consisting of public organizations, companies and independent users looking for open source solutions for maintaining their thesauri, code lists and authority resources. The focus on collaboration, the differentiation of user roles and the workflow management for content validation and publication have been the strengths of the platform, especially for those organizations requiring a distributed, yet centrally controlled, publication environment. In 2017, a new, completely reengineered, version of the system has been released, broadening the scope of the platform: funded by the ISA2 programme of the European Commission, VocBench 3 offers a general-purpose collaborative environment for development of any kind of RDF dataset (with dedicated facilities for ontologies, thesauri and lexicons), improving the editing capabilities of its predecessor, while still maintaining the peculiar aspects that determined its success. In this article, we review the requirements and the new objectives set for version 3, and then introduce the new characteristics that were implemented for this new incarnation of the platform

Keywords: Collaborative Editing, Ontologies, Thesauri, Lexicons, OWL, SKOS, OntoLex

1. Introduction

In 2008 the group for Agriculture Information Management Standards (AIMS) of the Food and Agriculture Organization of the United Nations[1] (FAO) developed a collaborative platform for collaboratively managing their Agrovoc thesaurus [1]. The so-called "Agrovoc Workbench" soon met the interest of other FAO departments and several other organizations interested in open source solutions for collaborative thesaurus development.

Rebranded as VocBench (VB), to suggest a more general environment for thesaurus management, the platform was later strongly reengineered in the context of a collaboration between FAO and the ART[2] group of the University of Rome Tor Vergata. The result of this collaboration, VocBench 2 (VB2) [2], released in 2013, was rethought as a fully-fledged collaborative platform for thesaurus management, freely available and open-sourced, offering native RDF support for SKOS [3] and SKOS-XL [4] knowledge organization systems [5], while retaining from its original version the focus on multilingualism, collaboration, and a structured content validation and publication workflow. Under the hood, the original VB1 backend for RDF was replaced with the RDF Management framework Semantic Turkey (ST) [6,7], already developed by the ART Group.

---

[1] http://www.fao.org/

[2] http://art.uniroma2.it

The strengths of the platform included the possibility for project administrators to define roles with very specific capabilities and to assign them to different users according to their proficiencies and authorizations, together with the publication workflow where dedicated users could supervise the work of others and accept their modifications. They were appreciated especially by those organizations requiring a collaborative yet centrally controlled publication environment. Unfortunately, this controlled approach was realized on a rigid model comprising a few first-class resources (e.g. concepts, concept schemes, etc.) and predefined operations on them. Consequently, some other users felt more freedom on data modeling was missing from the system, desiring unrestricted capabilities for editing data at its very core, as in triple-oriented RDF editing environments. The increased freedom would also reflect in the possibility to customize the models, going beyond plain SKOS/SKOS-XL modeling, which is sometimes a must for users dealing with complex, yet still KOS-like, resources.

VocBench 3 (or, simply, VB3) was planned to overcome the above limitations, while broadening the original scope of the platform to a general-purpose collaborative environment for development of SKOS thesauri, OWL ontologies and RDF datasets in general. The VocBench 3 project is funded by Action 1.1 of the ISA[2] Programme of the European Commission for "Interoperability solutions for public administrations, businesses and citizens"[3]. The action is managed by the Publications Office of the European Union[4]. VB3 has been developed in close collaboration with the ART group of the University of Rome Tor Vergata, the same group that contributed to the development of the second version of the platform.

VocBench 3 was released to the public on September 2017, under a BSD 3-clause license[5]. Since then, a second development iteration was carried on and terminated on July 2018 with VB3 v.4.0[6] (while v.2.0 and v.3.0 were released in the course of the iteration). The VocBench site[7] contains documentation, download links and other references. A third development iteration is being carried on, started after the fourth release of VB3 and terminating on June 2019.

In this article, we review the original requirements that drove the development of the platform and introduce the new objectives set for VB3. We then describe

and discuss the new features and architectural improvements that have been implemented to meet the goals for this next iteration of the platform. Our aim is thus to highlight the improvements over VB2, while we refer the reader to [2] for a general introduction to the system and for a comparative analysis of VB with related works. This article is a revised and expanded version of a previous work [8], in which we previewed VB3 just before the completion of its first development iteration. Improvements over that work include:

- discussion of related systems,
- discussion of additional features (e.g. extended input/output, integration of collaboration platforms),
- a more thoroughly analysis of the (meanwhile improved) support for OntoLex-Lemon,
- an example about XKOS showing the ability to support extensions to the core SKOS model,
- description of the system architecture,
- impact assessment.

2. Requirements

We list here the requirements that drove the development of VB3, including those originally put together for the development of the original VocBench platform and of its successor VocBench 2 (R1-R7 in the list below). The original requirements have been reassessed and reformulated, accounting for the widened scope of the platform and its improved editing capabilities, while new ones (R8-R15) have been added to complete the target objectives for the platform. The set of revisions and new requirements has been collected following proposals by the VocBench developing team and requests performed by stakeholders through several means: dedicated stakeholder meetings held at the Publications Office, ISA[2] programme, pro-active feedback on the support mailing lists of the platform. The input from these different sources all contributed to lay down the program for the development of VocBench 3, revised by the Publications Office and finally validated and approved by the ISA[2] committee.

R1. Multilingualism. Properly characterizing knowledge resources in different (natural) languages is fundamental. This especially holds for thesauri, due to their use in information retrieval, though the overall

---

importance of elaborated lexicalizations is progressively gaining momentum, thanks to data publication initiatives such as the Linguistic Linked Open Data[8] (LLOD) and to models for Ontology-Lexicon interfaces, such as *lemon* (lexicon model for ontologies) [9] and its most recent specification OntoLex-Lemon [10], realized by the eponymous W3C community group[9].

R2. Controlled Collaboration. One of the keystones of the system is to enable collaboration on a large scale: several, distributed users have to be able to collaborate remotely on a same project. Opening up to communities is important, though the development of authoritative resources demands for the presence of some control to be exerted over the resource lifecycle: for this reason, users must be granted different access levels, and some of them should be allowed to validate other users' work before it is committed to the dataset.

R3. Data Interoperability and Integrity. Interoperability critically depends on data integrity and conformance to representation standards. However, flexible models such as SKOS translate to underspecified possibilities on the one hand, and formal constraints beyond the expressiveness of OWL on the other one. Additionally, the increase in the offer of – often overlapping – standards in the family of RDF languages resulted in the necessity for systems to be flexible enough to properly read and manage all of them. It is thus important that VocBench enforces a consistent use of these models, by preventing the editors from generating invalid data, and by providing "fixing facilities" for spurious data acquired from external sources. Finally, support for alignment to other datasets is also an interoperability must for the Linked Data World.

R4. Software Interoperability/Extensibility. The system should be able to interact with (possibly interchangeable) standard technologies in the RDF/Linked Data world, with the possibility to surf linked open data on the Web, accessing SPARQL endpoints, resolving RDF descriptions through HTTP URIs, etc. as well to import/export data through standard Graph Store APIs and the like. The system should support extensions (sometimes called plugins), which can provide additional capabilities, often bound to predefined extension points: for example, enable to export data to a new type of destination (e.g. an SFTP server), or introduce a further serialization format (e.g. Zthes[10]).

R5. Data Scalability. The system must deal with (relatively) large amount of data, while still offering a friendly environment. This holds for some thesauri as well as for most lexicons. The user interface should appropriately subdivide data loading into subsequent requests and implement solutions for large results.

R6. Under-the-hood data access/modification. While a friendly user interface for content managers/domain experts is important, knowledge engineers need to access raw data beyond the usual front ends, as well as to benefit from mass editing/refactoring.

R7. Adaptive Context and Ease-of-use. Expanding a requirement on its predecessor, VB3 should provide an even smoother experience, with very low installation requirements and an as-short-as-possible time-to-use. Whether (and proportionally if) the user is an administrator configuring the system, a project manager configuring a project, a user requesting registration and connection to a given project, or a new user willing to test the system as a desktop tool without settings and configuration hassle, VB3 should adapt to their needs.

R8. RDF Languages Support. In contrast to its predecessors, dealing with thesauri only, VB3 has to support SKOS (/SKOS-XL) thesauri, OWL ontologies, and RDF datasets in general. Support for OntoLex-Lemon lexicons and for ontology-lexicon interfaces was introduced later as a further requirement [11,12].

R9. Maintainability (Architecture and Code Scalability). In particular, the ability to meet new requirements, cope with changed environments and make future maintenance easier. A weak spot of VB2, VB3 aims to achieve high levels of architecture/code scalability. In VB3 it is mandatory to be able to add new services, functionalities, plugins, etc. without the fabric of the system being altered or too much effort being required to align these new elements with all the characteristics of the system, such as validation, history management, roles and capabilities.

R10. Full Editing Capability (RDF Observability and Reachability). Any complex RDF construct should always be inspectable and modifiable by users (providing they have the proper authorization) even in its finer details. While the platform can provide high-level operations for conveniently creating/modifying complex descriptions of resources according to predefined modeling design patterns, the user should never be prevented from inspecting/altering these elements.

R11. Provenance. Actions in VB3 should be handled as first-class citizens themselves, being identified and qualified by proper metadata, logged in a history with information about which user performed an action,

---

when they did it, which parameters have influenced its performance, etc. Metadata answering to the five "Ws" (with the possible exception of the "why") should provide all information for tracking the origin of an action.

R12. Versioning Support. Besides history and validation, providing triple-grained information about actions enriched with provenance metadata, it should be possible to take static, periodic, snapshots of datasets.

R13. Dataset-level Metadata Descriptions. For the Semantic Web to fully achieve its vision, linked open data has to speak about itself [13]. This means not only having data modeled according to well-known shared vocabularies, but to be able to grasp meaningful information about a dataset without having to dig into its content. Edited datasets should be coupled with resuming information about their characteristics, that can be published together with them.

R14. Customizable User Interface. User interfaces merely based on ontology description are limited to the analysis of the axiomatic description of the resources they show and of their types, ignoring possible desiderata of the user. VB3 should allow users to represent the information that they want to specify at resource creation, per resource type, so that it will be prompted to the user. Connected technical aspects, such as proper transformation of the user input into serializable RDF content, should also be covered.

R15. Everything's RDF. VB2 used a relational database to store user and project management information as well as history and validation information. Conversely, VB3 should follow a more uniform approach, adopting RDF for virtually any storage need.

3. Architecture

VB3 is based on a classical three-tier architecture (Figure 1), structured through a presentation layer, a service layer and a data layer. The web application (developed using Angular[11]) is a front end for Semantic Turkey, which, with respect to its former version adopted in VB2, has then evolved into a fully-fledged collaborative environment for RDF management.

3.1. A Lesson Learned: Standards Compliancy Does Not Guarantee Interchangeability

VB3 does not feature the RDF abstraction layer – a middleware allowing for different RDF technologies

to be interchanged in the lower tier – that characterized its earlier incarnations. This neat drift in the architecture is due to the acknowledgement that, at least with the current technologies and specifications, it is utopic to expect two API implementations to be practically swapped under the same functional layers. While the first VocBench largely adopted triple-oriented API to access the data and the second introduced SPARQL queries, VB3 is completely based on SPARQL, which is the weak point where technology interchangeability fails. By first, standards compliancy is not to be taken for granted when non-common characteristics of the language (e.g. bindings) are being used, as different triple stores might miss to implement (or partially implement) them. Second, and most important, the SPARQL specifications do not constrain the order in which the various clauses of a query should be resolved: the strategies of SPARQL query resolvers and optimizers may vary dramatically across diverse triple stores, making it difficult to express a certain information need as a SPARQL query that will behave as expected under all conditions. Additionally, there are critical differences in the way stores organize content (e.g. where the inferred triples are stored and how they can be retrieved). Fighting these idiosyncrasies is not an easy war, especially in an editor that must guarantee flexibility and that cannot perform any content-oriented optimization (as the datasets it manages are not know a priori).

With the advent of RDF4J[12] under the stewardship of the Eclipse Foundation, we decided time was mature for taking a stance and adopt that RDF framework as the official API for ST. Notably, we embraced the RDF4J's sail-mechanism (Storage and Inference Layer) for developing storage-side components (represented under the local and remote triple stores in the architecture view) for intercepting effectively-modified triples (see section 3.2 and section 4.3).

3.2. Extendible Architecture

The adoption of OSGi (formerly Open Service Gateway initiative)[13] allows for dynamic plugging of extensions (see requirements R4 and R9). We have thus developed a very rich mechanism for describing extensions: there is a general concept of *component*, an object that can be identified in the system, configured, and scoped (to four domains, i.e. the whole SYSTEM, a USER, a PROJECT or spaces defined by
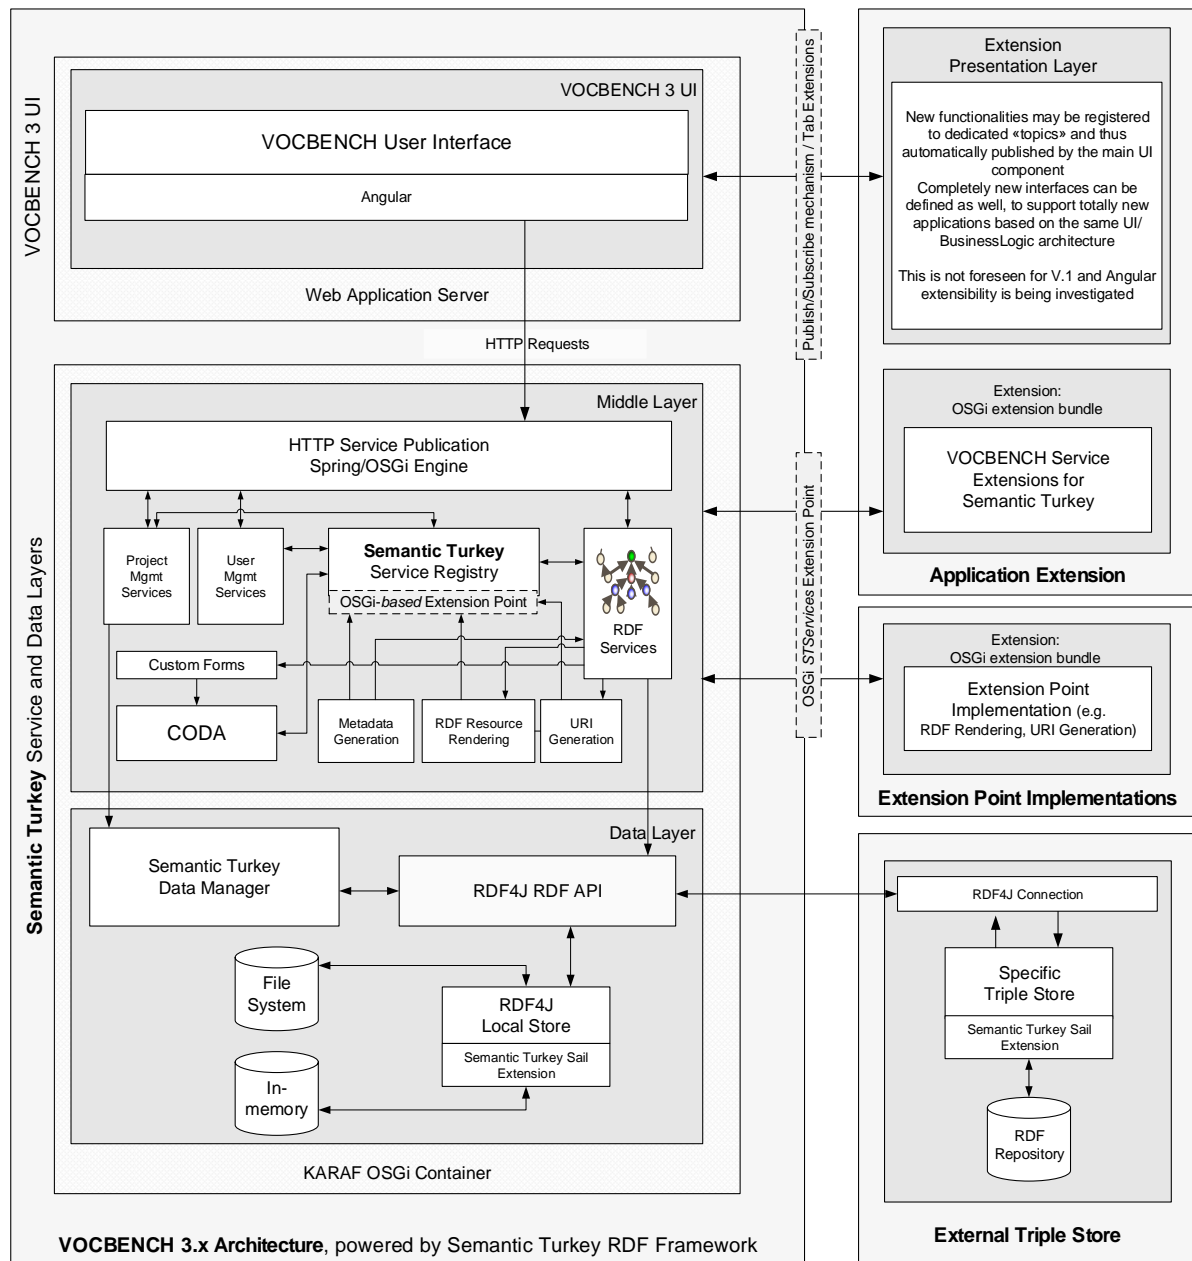
Figure 1. VocBench 3 Architecture

<USER,PROJECT> pairs). *Extension points* are components defining extensible functionalities, such as data loading/deployment, rendering of RDF resources on the UI, connection to and interaction with collaboration platforms, etc. Developers willing to extend VB3 with respect to a functionality can supply an implementation of its associated extension point.

While there is currently no native extension mechanism for web interfaces built through Angular, the user interface of VB is informed by the extension point mechanism and supports extensions with automatically built forms for configuration of settings, selectors for choosing the implementation to choose, etc.

Currently, numerous extension points have been included in the system, in order to make extensible the

user-visible features of VB3 (see Section 4). Hereafter, we will describe three extension points that are not covered elsewhere in paper, because they are related to lower level aspects of the system.

- *RepositoryImplConfigurer*: this consists in configuration templates for different triple stores. Currently, there are implementations for the storage solutions of RDF4J and for Ontotext GraphDB[14] (formerly OWLIM [14]). As anticipated before, these extensions support configuration dialogs (see Figure 2) including options that are specific to the chosen storage solution. Figure 2 shows the options for the RDF4J Native Store: whether to force the syncronization of the filesystem to the non-volatile storage device (a performance vs durability tradeoff), which indexes to use (depending on the query patterns), and which types of reasoning perform (which are not costless).
- *SearchStrategy*: a search component that can use non-standard, triple store specific text-search capabilities. Indeed, SPARQL only supports FILTER with some string matching functions. However, as FILTERs are often applied to discard solutions rather than find them, their use as a search mechanism is quite inefficient if the number of candidates is very high (e.g. tens of thousands). In contrast, a dedicated *SearchStrategy* can take advantage of the efficient fulltext capabilities provided by GraphDB.
- *URIGenerator*: a pluggable component for the automatic generation of URIs. Indeed, there are many strategies to mint URIs, based on technical and political arguments. Ontologies often use human-friendly but language-specific identifiers based on a label (e.g. `http://schema.org/Person`), to make data more easily consumable (without specific tool support). Conversely, thesauri and other multilingual datasets frequently adopt alphanumeric, language independent, identifiers (e.g. `http://aims.fao.org/aos/agrovoc/c_2993`). It is possible to develop new generators for custom patterns that are unsupported by the (highly configurable) *URIGenerator* packed with VB.

3.3. Project and User Management

As anticipated in the description of the architecture, *project* and *user management* have been completely

Figure 2. Dynamic dialog for the configuration of a repository based on some store solution (in this case, the RDF4J Native Store)

rewritten as part of the Semantic Turkey framework. The system offers an abstract representation of the core entities managed: users, projects, system properties, etc. and API for storing/retrieving them, so that different implementations can be provided.

The default implementation represents another aspect which is streamlined with respect to VB3's predecessors: it is completely file-based. Consequently, the relational DB previously used for system administration (and for metadata representation, which is now completely represented in RDF, see section 4.3) is no more necessary. The use of the filesystem is not just a choice to get rid of the relational DB. An accurate organization in the distribution of the descriptors for users, projects, configurations, settings, plugins (these last three can in turn be also scoped differently) and their relationships enables an easy porting of data across different distributions: it is thus possible to easily transfer all data to another installation of Semantic Turkey, or to separately move users, projects, deciding whether to copy or not their settings etc.

3.4. CODA

CODA[15] [15] is a is an architecture and an associated Java framework for the triplification of results from analysis of unstructured content. The facilities provided by CODA include a powerful language – PEARL[16] [16] – for projection and transformation of annotated content into RDF. While its most natural and implied use – acquisition of knowledge from text – has not yet been exploited inside VB3, CODA has already found diverse applications, including *custom*

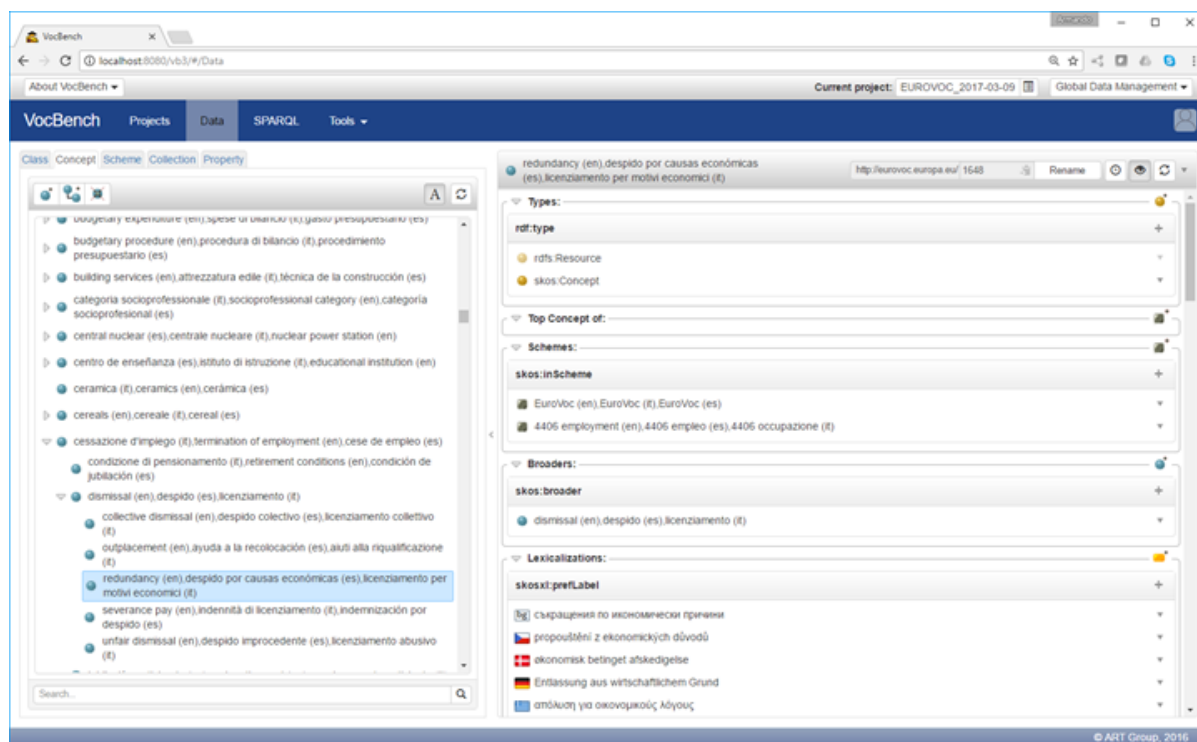Figure 3. VocBench UI showing EuroVoc (`http://eurovoc.europa.eu/`)

forms (see section 4.1.5) and Sheet2RDF[17] [17], a platform (being integrated into VB3) for the acquisition and transformation of spreadsheets into RDF.

4. VocBench 3 New and Improved Features

In this section we list features, functionalities and visual changes that are more evident to the user than the description provided in the architecture. Conformance to and satisfaction of the requirements expressed in section 2 is reported case by case.

4.1. User Interface (UI).

The overall data view (Figure 3) preserved its overall organization, with data browsing views on the left and the description of the selected resource on the right. However, there are notable changes on both sides.

4.1.1. Browsing Data Structures

VB only provided a concept tree to browse the content of the dataset, supposed to be a SKOS(-XL) thesaurus. In line with its wider scope, VB3 now offers

different data structures in different tabs, depending on the modeling vocabulary. OWL offers three tabs with a class tree and instance list, a property tree and a datatype list respectively, while SKOS adds to them a concept tree, a list of schemes and a collection tree (showing the containment between collections). Onto-Lex-Lemon adds tabs for lexicons and lexical entries.

4.1.2. The resource-view

The several tabs of VB2 (which inherited and extended the tab-based model of VocBench 1) that populated the "concept details" panel have been replaced with a single component, called *resource-view*.

In contrast to both VB1 and VB2, there are no first-class citizen resources, such as SKOS concepts and SKOS-XL labels; in fact, all resources now can be viewed and edited through the *resource-view*. This capability supports generic RDF development (requirement R8) and marks a profound departure from previous versions of the system, which had dedicated panels for certain types of resources that occur in thesauri (e.g. concepts and reified labels). The general applicability of the resource-view and the possibility to edit any of their details also contribute to satisfy requirement R10.

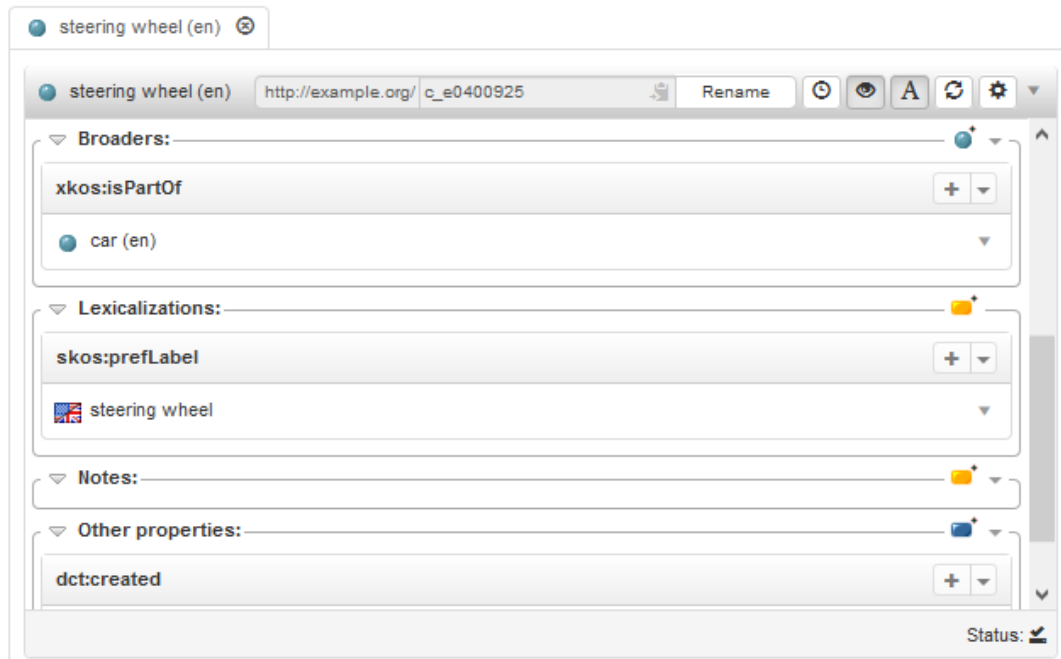[17] `http://art.uniroma2.it/sheet2rdf/`

Figure 4. Visualization of specialized hierarchical relationships between SKOS concepts

The resource-view adapts to the inspected resource: a few sections are shared among all resources, such as *types*, listing the rdf:types for the resource, *lexicalizations*, listing the available lexicalizations and *properties*, listing properties not addressed by other sections, while others are specific to the inspected resource. While the mapping of these sections to properties of the core modeling vocabularies is trivial (e.g. *types* to rdf:type), the sections are however presented with a predicate-object style in order to qualify the predicate, as they might include user-defined or domain-specific subproperties of the above ones. For example, see Figure 4, instead of saying generically that "car" is "broader" than "steering wheel", we can we can use XKOS [18] to state that the latter is a meronym of the former. Indeed, a subproperty of skos:broader can be indicated (optionally) both in the dialog for the creation of a new concept (when creating it as narrower of an existing concept, thus refining the relation among them) and in the dialog for the addition of a broader concept. The previous example shows the ability of the user interface to handle extensions (e.g. XKOS) of the core modeling vocabularies (e.g. SKOS), contributing to fulfill requirement R8.

4.1.3. Dataset Model and Lexicalization Model

The *lexicalizations* section abstracts different properties and modeling patterns for lexicalizations, as it offers specific resolution of their shape, always showing the form of the lexicalization, i.e. merely the label for rdfs:label and for SKOS terminological labels, the skosxl:literalForm for SKOS-XL labels and the ontolex:writtenRep of the canonical form associated with the lexical entry lexicalizing the resource in OntoLex-Lemon. In VB3, the concept of *lexical model* has been introduced (and separated from the *knowledge model*, e.g. OWL or SKOS) so that, for instance, it is possible to select SKOS-XL as a lexical model for both OWL ontologies and SKOS thesauri. In OntoLex-Lemon, the indirection is perhaps the most evident as one possible path from the lexicalized resource to the shown lexicalization is:

ontolex:isReferenceOf → <ontolex:Sense> → ontolex:isSenseOf → <ontolex:LexicalEntry> → ontolex:canonicalForm → <ontolex:Form> → ontolex:writtenRep → *<shown lexicalization>*

Many other paths are considered, due to shortcuts on property chains and inverse properties. So, the written representation is shown, but the user can click on it and inspect the full series of linked RDF terms.

This revised model greatly improves requirement R1 by covering not only diverse natural languages, but the different formal languages (and thus, R8 as well) in which the lexical information can be encoded.
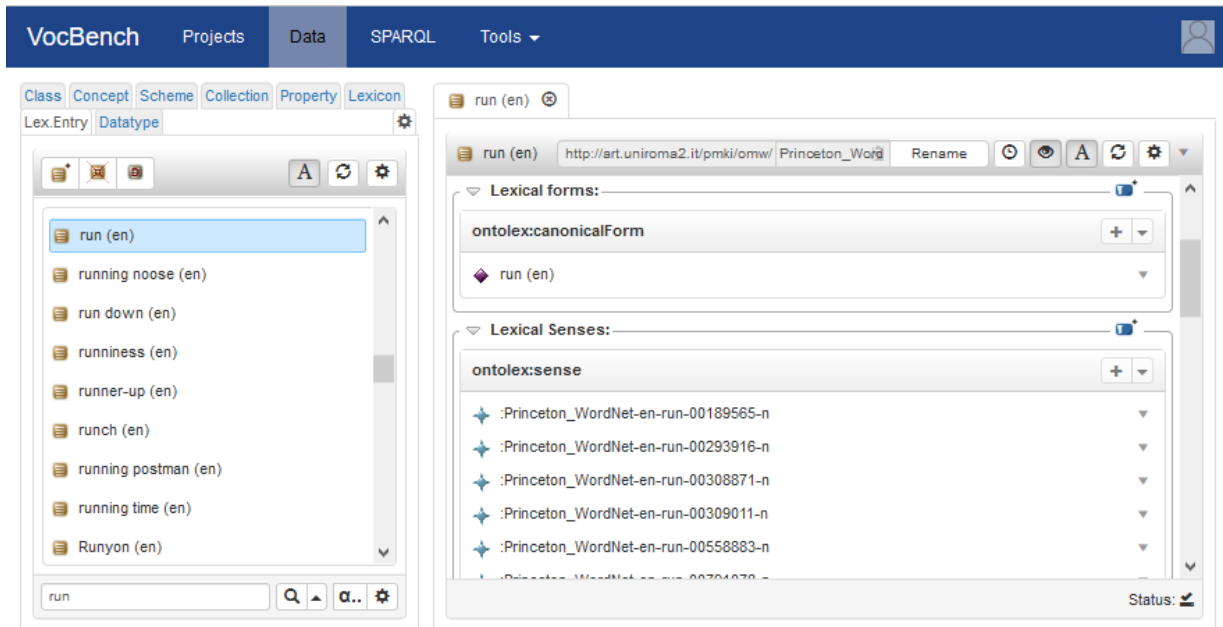
Figure 5. Visualization of WordNet's Lexical Entries in OntoLex views. In fact, the system is managing the whole collection of 34 wordnets collected by Open Multilingual Wordnet [61]

4.1.4. Support for Lexicons and Ontology-Lexicon Interfaces

The support for lexicons required, even more than with thesauri, a data-scalable user interface (see Figure 5). Modern thesauri are (usually) organized around a hierarchy of concepts that can be easily browsed by progressively expanding the explored branches of the tree. Conversely, lexicons have huge flat lists of entries that are difficult to represent. Additionally, the OntoLex-Lemon model introduces the notion of onto-lex:LexicalConcept, corresponding to the notion of *synset* in lexical databases following the model of WordNet[18] [19]. In wordnets, the set of synsets is arranged in a structured tree for what concerns synsets related to *nouns*, but offers a shallow hierarchy (mostly, a horizontal list of elements) for other parts-of-speech, especially for what concerns *verbs*. We have thus offered different browsing modalities, based on the exploration of the full content, organized according to different data structures (e.g. trees for concepts, classes and properties, lists for schemes and lexicons or single/double-character indexed lists for lexical entries) or a search-based exploration, that uses the search functionality to selectively show matching entries in their associated panel, thus guaranteeing scalable solutions (req. R5) depending on the nature and specific size of each loaded dataset.

VB1 and 2 showed concepts through their labels in all the selected languages for visualization. In VB3, an option allows for toggling between the URIs/qnames[19] (qualified name) of the resources and the string composed by an implementation of the extension point *rendering engine* (see Section 3.2). The default renderer behaves like VB1 and VB2, showing labels in all the selected languages for visualization (again, R1), being configurable in the languages to show.

4.1.5. Custom Forms

An important new aspect of the user interface is offered by *custom forms*, a flexible data-driven form definition mechanism that we devised for VB, allowing users to perform a declarative specification of the key elements that concur to the creation of a complex RDF resource (satisfying req. R14).

Custom forms have been described more in details in [20], which analyzed and evaluated their expressive power by applying them to the use case of representing

---

Figure 6. Custom Form for a relational noun in the OntoLex-Lemon model

lexical entries using the OntoLex-Lemon vocabulary. In that work, a subset of the *lemon* Design Pattern Library [21] was implemented as custom forms: VB exploits them to generate a form-based interface for the creation as well as the visualization of diverse lexical entries. Figure 6 reports a form filled with information regarding the entry "director", which is said to denote the property dbo:director in the DBpedia ontology. Furthermore, the syntax-semantics interface is defined by establishing the correspondence between the atom x dbo:director y and the stereotypical verbalization *y is the director of x*.

4.2. Controlled Collaborative Editing through Role-based Access Control (RBAC)

A single installation of VB can handle multiple projects, which can also be interlinked for mutual data access (e.g. for purpose of alignment). VB promotes the separation of responsibilities through a role-based access control mechanism, checking user privileges for requested functionalities through the role they assume (req. R2). While VB2 had hard-wired and scarcely configurable roles, which do not easily scale-up to possible extensions of the system (req. R9), in VB3 we have created a dedicated language for specifying capabilities in terms of *area*, *subjects* and *scopes*. E.g. the expression:

auth(rdf(datatypeProperty, taxonomy), 'R')

corresponds to the requested authorization for being able to read taxonomical information about datatype properties. The 'R' stands for READ (or, equivalently, RETRIEVE), as in the CRUD paradigm, rdf is the *area* of the requested capability while datatypeProperty and taxonomy define the *subject* and *scope* respectively of the capability. We recall that "CRUD is [...] the most common acronym for the four basic functions of persistent storage: create, retrieve, update and delete" [22]. Our language supports all these functions, which are identified by their initial letter. Actually, we extended CRUD with a fifth letter 'V' standing for VALIDATE: this permission grants the right to validate other users' work (see Section 4.4).

The authorization policy is implemented as a series of facts in Prolog [23], whose resolution mechanism provides a principled mechanism for authorization decisions. Indeed, the same set of expressions is manipulated and validated on both the server and the client, by using different technologies, respectively, tuProlog[20] [24] and jsprolog[21]. While the ultimate authorization is performed by the server, the client does a similar evaluation to show/activate UI elements depending on the permissions of the logged user. This way, it is possible to have very dynamic UIs automat-

---

[20] http://apice.unibo.it/xwiki/bin/view/Tuprolog/WebHome
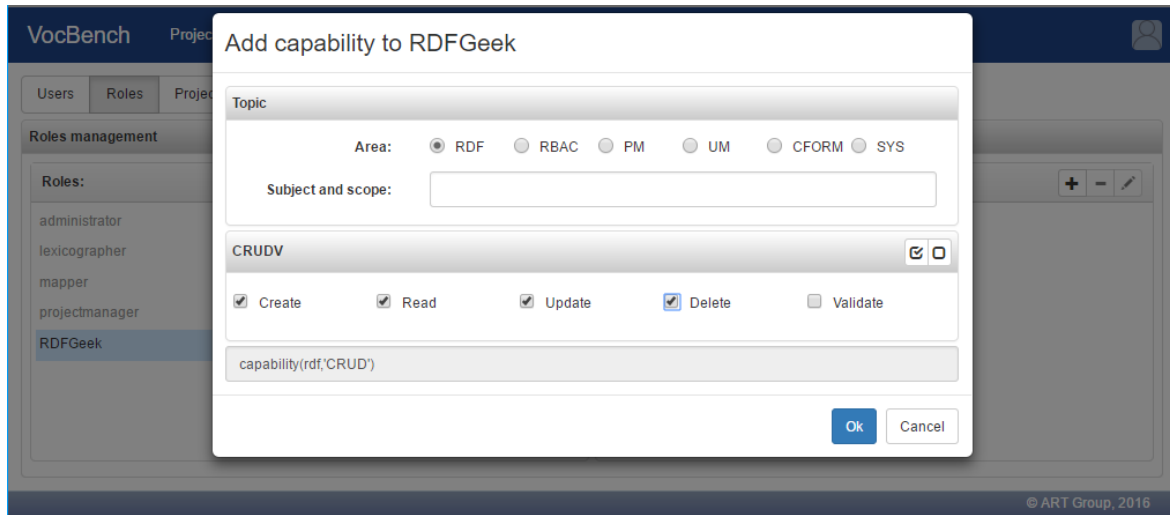
[21] https://github.com/Sleepyowl/jsprolog/

Figure 7. Editing a capability for a role in VocBench

ically modeled on users' capabilities, which are ultimately defined in the services, with no redundancy in the code.

New roles can be easily created, and existing ones can be modified, through a dedicated *rbac* editing wizard (Figure 7). The default policy recognizes typical roles and their acknowledged responsibilities:

- *Administrator*: the sole inter-project role (i.e. the role exists a-priori from projects). The administrator has, by definition, full access to the system.
- *Project Managers*: project-local administrators who can do everything within the boundaries of the project(s) they have been assigned to with that role.
- Specific project-local roles: *ontology editors* (allowed to make changes at the axiomatic level), *thesaurus editors* (allowed to work on thesauri without OWL editing actions), *terminologists/lexicographers* (allowed to edit lexicalizations, can be limited to certain languages), *mappers* (allowed to perform alignments only), *validators* (allowed to validate others' actions, see Section 4.4) and finally *lurkers*, that can read everything in a project but have no editing authorization

4.3. Advanced History and Change Tracking mechanism

In VB2, the *change tracking* mechanism that powered *history* and *validation* was based on a predefined set of recognized operations, severely limiting maintainability (req. R9) and the possibility to perform (req.

R6) under-the-hood changes (e.g. through SPARQL) *while* keeping a complete history of the dataset. In VB3, we implemented a completely new track-change mechanism working at triple-level. For any action, triple additions and removals are intercepted and stored into a separate RDF repository (the *support repository*) (req. R15) together with metadata about the action (req. R11). This design was guided by an analysis [25,26] in which we discussed the nature and the representation of change, reviewed relevant version control systems for RDF, and delved into the challenges posed by validation.

4.4. More Powerful yet Streamlined Workflow Management

VB2 had a 5-step publication workflow, clocked by the property "status" (with values: proposed, validated, published, deprecated and proposed_deprecated) and, redundantly, with information stored in the DB about the status of operations to be validated. Also, in VB2, the concepts of *resource* and *action* were mixed up in the validation procedure, with the status of a resource being affected by the validation (e.g. moving from "proposed" to "validated"), while single changed triples had no trace of their validation status.

This follows from the fact that it is not possible to attach a status to a triple in RDF, if not by reifying the triple. An in-depth discussion of different reification mechanisms can be found in [27]. Some mechanisms deeply affect how data is represented (e.g. singleton properties [28], or even standard reification), while named graphs [29] would be sufficiently lightweight,
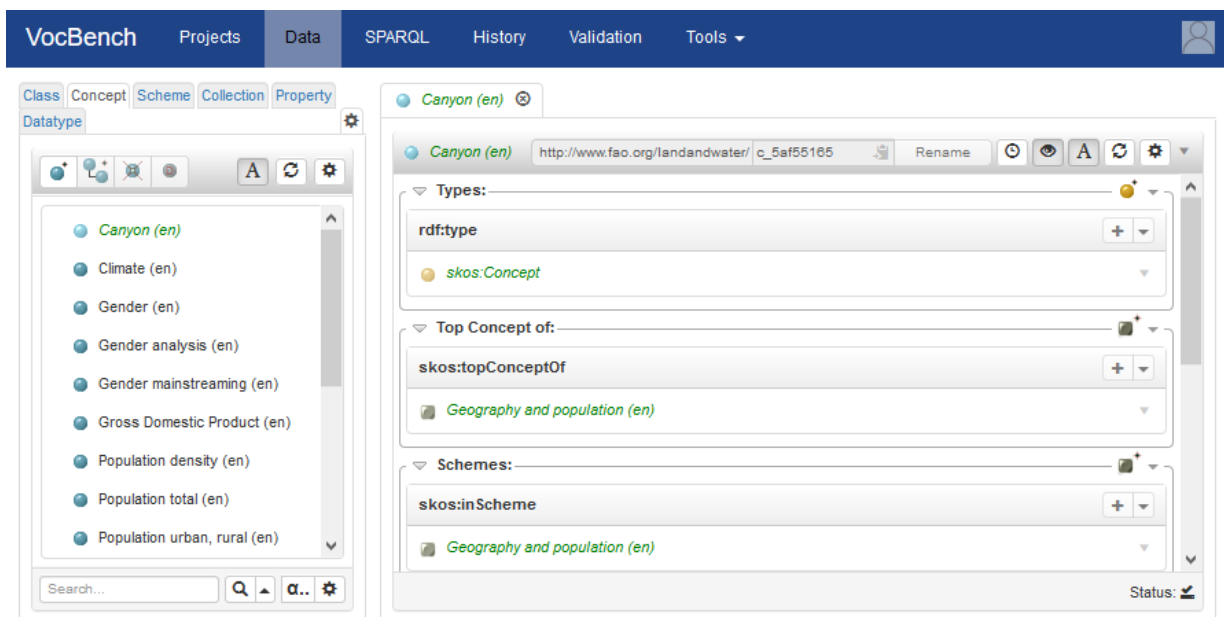
Figure 8. Proposed concept "Canyon" shown in the concept tree, as well as in the resource-view

if VB did not already use them to differentiate between local and imported triples. Nonetheless, as discussed below, VB3 eventually adopted named graphs for the representation of proposed triples.

Benefiting from the new change tracking system, we have made things clearer and easier: there is no "status" property anymore, as the workflow is implicitly expressed by the validation mechanism coded into graphs. The added/removed triples are stored in the support repository as described in the previous section, while non-reified "previews" of them are available in separate graphs (*staging-add-graph* and *staging-delete-graph*) in the main repository, while the system presents them appropriately to the user. In this way, the *main graph* (i.e. the graph where managed information is stored) represents stable information, which does not need to be tagged as "validated". The distinction between "validated" and "published" has been removed as it has never been put in place in any known user workflow. Finally, the status of deprecation has been represented through the official owl:deprecated property. The status of "proposed deprecated" is also intuitively represented by the need to validate the action for setting the owl:deprecated property to "true".

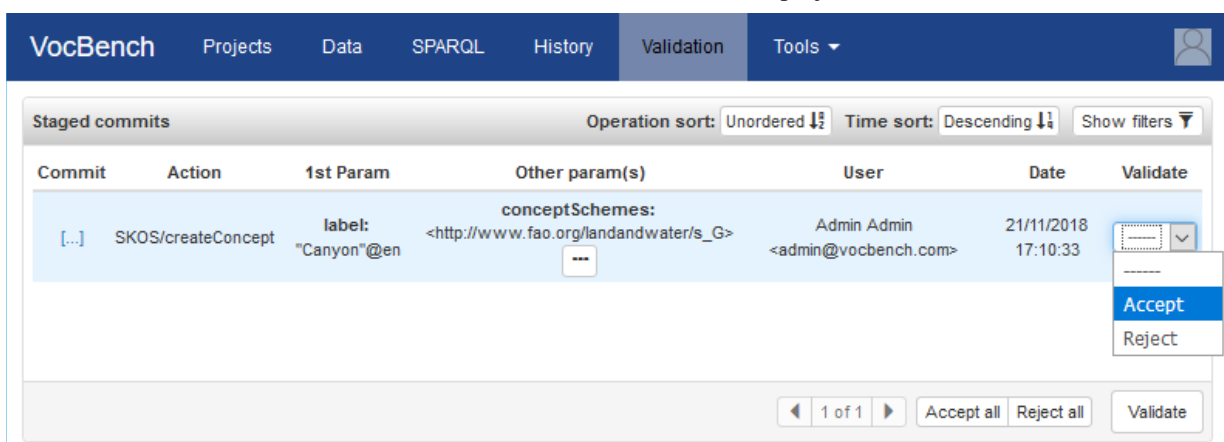In Figure 8, we show the concept "Canyon" that was created in a project with validation.

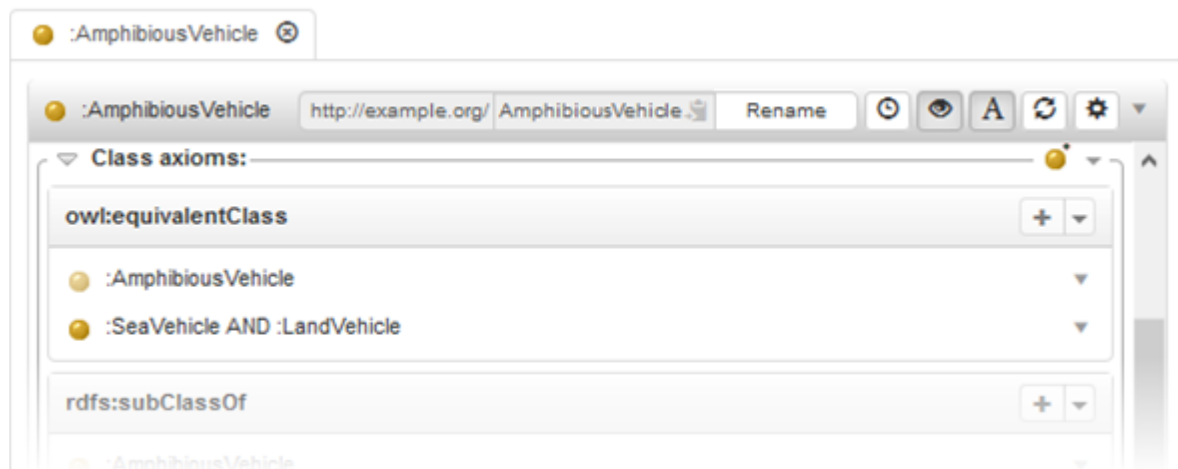Figure 9. Accepting the creation of a concept to move it from "proposed" to "validated"

Figure 10. Resource-view of the class `AmphibiousVehicle`, showing the use of the Manchester syntax to visualize class descriptions and the different rendering of inferred statements (lighter color). The button with the eye icon controls the inclusion of inferred statements

The triples describing this concept are asserted in the *staging-add-graph* (instead of the project *main graph*): the resource-view shows them with a combination of green color and italic font, while triples in the *main graph* (none in this example) are displayed in black with a normal font. Since the assertion of the characterizing type (`skos:Concept`) is also staged for addition, the concept itself is considered proposed, and thus displayed in the concept tree differently from other (already validated) concepts. Figure 9 shows the creation of the concept in the list of operations pending for validation: a user with suitable rights (see Section 4.2) can decide to either accept or reject these operations. Accepting an operation makes its effects permanent (i.e. updating the *main graph*), while rejecting an operation undoes its effects (i.e. remove the previewed triples in the *staging* graphs). In the latter case, no trace of the operation remains, nor is the operation logged in the history: rejecting an operation should be like the operation was never performed. Future directions for the system foresee the possibility to store (in a logically-separated space) rejected actions so that, should they be re-proposed in the future, the user can be informed and discouraged from submitting them for validation again.

4.5. Improved and More Complete Support for SKOS

VB2 had already an advanced support for multiple SKOS schemes. We have improved the management by allowing users to select more schemes for browsing the concept tree and by adopting a combination of conventions and editing capabilities for quickly associating the proper schemes to newly created concepts and collections. Support for SKOS collections and ordered collections has been introduced in the system with dedicated UI views and editing facilities.

4.6. OWL Support

VB already allowed for importing ontology vocabularies for modeling thesauri. Now VB also supports ontology development (requirement R8) with an almost complete coverage of OWL2 and using the Manchester syntax for both editing and visualization of class expressions (see Figure 10). VocBench delegates reasoning to the triple store that manages the ontology and its data (see Section 3). Nonetheless, VB differentiates between explicit and inferred statements: they are rendered differently in the user interface, and users may toggle their inclusion in the resource-view (see Figure 10). Currently, VB does not present justifications of the inferences, because RDF4J (see Section 3.1) lacks any mechanism to obtain them.

Triple stores usually do rule-based reasoning with a total materialization strategy: the entailment closure of the ontology is computed in advance by the iterative application of the rules defining the semantics of the ontology modeling language. In the rest of the section, we provide more details about the triple stores commonly used with VB highlighting differences in flexibility, expressiveness and suitability to workloads.

RDF4J used to ship a reasoner with hard-coded rules for RDFS. Moreover, when knowledge is retracted from the ontology, this reasoner computes the entailment closure from scratch, since it cannot determine which entailments are no longer supported. This

Figure 11. Export data

approach is clearly inefficient when the workload includes many deletions that affect the entailment closure slightly, e.g. the deletion of a fact assertion is usually less impactful than a deletion at the schema level. An alternative "schema caching" reasoner is not rule-based, since it gathers schema-level assertions and creates a data structure to quickly determine inferred statements. This reasoner better handles deletions that do not affect the schema, while being more performant (up to 80x faster) and more complete (but still bound to RDFS). For these reasons, the forward-chaining reasoner was superseded by this new one in RDF4J 2.5.

Ontotext GraphDB implements reasoning using a rule-engine that can execute arbitrary rules written in a proprietary rule-language. GraphDB provides optimized rulesets for RDFS, OWL2-RL and OWL2-QL semantics. When statements are retracted, GraphDB first determines (by forward-chaining) their logical consequences, and then removes only those inferred statements for which it is not possible to determine (by backward-chaining) any derivation not including the knowledge being retracted. GraphDB can also perform consistency checking and prevent modifications that would produce inconsistent knowledge. GraphDB features a non-rule implementation of `owl:sameAs` that eliminates the need to store N^2 identity links and to rewrite statements for each equivalence class.

4.7. Input/Output

VB3 features completely redesigned data load and export capabilities. Since they are somehow specular, we will focus hereafter on the export procedure, using the example in Figure 11 dealing with export of a SKOS-XL thesaurus.

In the example, we accepted the default option to export the *main graph* (the data edited by the users).

In some circumstances, the triples in the chosen graphs can't be exported as they are. VB2 partially satisfied this need with the ability to export a cut of a thesaurus. VB3 clearly required a more general solution: the extension point *rdf transformer* (see Section 3.2) has been added, and data being exported can be processed by a chain of *rdf transformers,* each performing a destructive transformation (on a temporary copy of the data. In our example, we turned SKOS-XL reified labels into plain SKOS labels.

The RDF data being exported, optionally processed by a chain of transformers can be placed to a *file*, a *triple store* or a *custom destination* (the last two options associated with the extension point *Deployer*, see Section 3.2). The first two options require an *RDF reformatting exporter*, the extension point (see Section
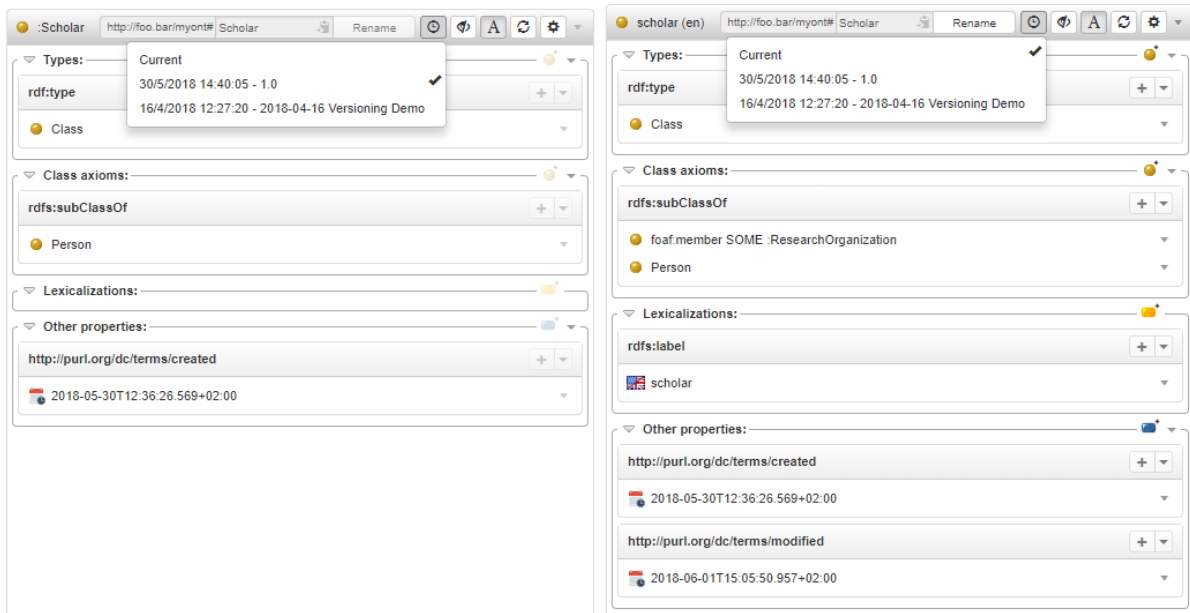
Figure 12. Time-traveling across different versions of a resource

3.2) for supplying serializers of the data into a byte sequence. In the example, we chose to deploy an RDF/XML serialization of the transformed data to an SFTP server.

It is worthy of note that these complex export chains can be saved and reused (even across projects). Furthermore, the configurations of each of the individual components (*transformers*, *reformatters* and *deployers*) can be saved and reused in different export chains.

4.8. SPARQL Querying and Update

A new SPARQL UI, based on YASGUI [30], has been included in VB3, featuring the same feeding-from-live-data mechanism present in VB2.

VB3 also introduced the possibility to store queries (at different scopes described in section 3.2) and share them with other users, even across projects. This capability is useful in different cases: i) a complex query can be reused, ii) a user writes a query for other users less proficient (or no proficient at all) with SPARQL, iii) periodic execution of a query (e.g. for analytics).

VB3 offers several options for downloading the results of a query, including various tabular formats for tuple queries, and RDF serializations for graph queries. Following the observation that the latter case is not much different from a data export, we supported a customizable chain of *transformers* as we did in the export. For instance, a general SPARQL query could extract all alignments available in a dataset, while a

chained RDF transformer could filter only those alignments referring to DBpedia [31]. The alignment export query can thus be reused across different cases.

4.9. Versioned datasets and metadata

In VB3, users can create snapshots of a repository (req. R12) and tag them with a version identifier (and other metadata, such as the time of creation of the snapshot). Users can travel across the different points in time identified by these versions, and thus analyze the evolution of browsed resources. The time travel can be performed both globally, by switching version so that everything in the UI refers to the selected version, and locally, by inspecting different versions of a resource in the resource-view (Figure 12) or different versions of a tree (of classes, concepts, etc.)

4.10. Alignment

VB3 provides the same inter-project alignment support of VB2, allowing users to browse other projects and supporting semi-automatic label-based searches over them to provide candidate resources for alignments. In addition to this on-the-fly generation of mappings, VB3 introduces a tool for loading and validating alignments following the model of the INRIA Alignment API [32]. Validated alignments can then be projected over standard RDFS/OWL or SKOS mapping properties, depending on the validated relation and the

Figure 13. Listing issues in JIRA associated to concepts in a thesaurus, opening one of these concepts and checking other collaboration options

involved entities. E.g. two classes mapped through an inria:EquivRelation will be mapped through the property owl:equivalentClass while two SKOS concepts will be proposed to be aligned with a skos:exactMatch or a skos:closeMatch.

4.11. Collaboration Environments

The extension point *collaboration backend* (see Section 3.2) has been added to the platform to enable the use of diverse collaboration environments with a predefined implementation for JIRA [22], the popular project management platform by Atlassian. Collaboration environments provide, in the context of defined projects, means to create and share tasks/issues/stories/discussions that can be assigned to users for their resolution. This enhances coverage of requirement R2 for a controlled collaboration. Figure 13 shows how the resource-view of a resource tells whether there are pending collaboration items for the resource and let users create new ones. VB3 also features a view over all pending items for a project. In several places, VB3 provides links to the page on the underlying collaboration platform for the details of a collaboration item.

4.12. Metadata Export

The "metrics" section of VB2 has been replaced with a page for editing and exporting metadata (R13)

modeled after several existing metadata vocabularies: the Data Catalog Vocabulary (DCAT) [33], the Asset Description Metadata Schema (ADMS) [34], The Vocabulary of Interlinked Datasets (VoID) [35] and the Linguistic Metadata vocabulary (LIME) [36] (a lexical extension to VoID). While DCAT and ADMS mostly deal with static metadata, VoID and LIME offer statistical information about the dataset and its lexical information. The information of VoID and LIME is being computed through the LIME API [37]. This metadata build and export functionality is implemented as an extension point of the platform, so that new vocabularies can be dynamically added to the platform through the extension point *Dataset Metadata Exporter* (see Section 3.2). For instance, an application profile for DCAT thought for European public sector data portals (DCAT-AP[23]) has later been added to the list of exporters, as of the ISA[2] context specifically supporting public administration. In the spirit of [38,39,40,41], a planned upgrade of the system will exploit conceptual and lexical metadata in order to automate the setup of alignment processes, identifying the key features of datasets that can be used by the matching systems and external resources that can be leveraged on support.

4.13. Integrity Constraint Validation

A section dedicated to Integrity Constraint Validation (ICV) allows the user to inspect possible anomalies. These include violations of formal constraints (e.g.

---

thesauri constraints on existence and uniqueness of preferred labels, disjointness between taxonomy and relatedness, etc.) or problematic (though not necessarily illegal) patterns (e.g. a skos:Concept having a broader concept and being the top concept of a same scheme). Interactive fixes are provided (req. R3) for each discovered integrity break.

4.14. Desktop Tool and Collaborative Web Platform

As of requirement R7, the system offers a very lightweight installation (i.e. unzip and click-to-run) which, followed by default configuration options for both system and project creation, makes VB3 a good choice for users looking for a simple and easy-to-use desktop tool. Other more complex settings are still possible, satisfying different needs for distributed installation (separation of data servers, UI servers), better performance, etc.

4.15. Declarative Service Implementation

The VB framework provides Java annotations for specifying diverse characteristics of a service: the need for read/write access to the data, the capabilities the user must have, the prerequisites on the input parameters (e.g. an RDF resource must be declared in the dataset), etc. Service development becomes easier, less prone to errors and the produced code is more readable, as developers can focus on what the service does.

5. Impact

A heterogeneous user community[24] has grown in these years around VB, including large organizations, companies needing VB as users and companies featuring it as their platform of choice in their range of offered RDF services, consultants needing a modeling environment, etc.

Some of these long-lasting users still adopt VB2, but there are compelling reasons for which they will migrate to VB3: i) VB2 is no longer supported in terms of updates ii) the feature set of VB3 subsumes the one

of VB2, iii) some features of VB2 have been substantially improved when migrated to VB3 (e.g. search, history/validation, etc.). Additionally, we should mention that VB3 has already relevant users[25]. The Publications Office of the EU (which is managing the development of VB3) has already adopted VB3 in production for their EuroVoc thesaurus and for collaboratively managing the Common Metadata Model [42] ontology, thus exploiting the widened support of VB3 for OWL ontologies. The Publications Office is also, since this October 2018, providing hosting and support to Directorate Generals (DGs) of the European Commission for managing their datasets. At the time of writing (only 4 months after the decision to provide this support) the production installation of VocBench 3 at the Publications Office is hosting 47 projects for 10 DGs, with each DG and the Office itself having more than one team working on different projects. The number of adopters exploiting this supported hosting within the European Commission and related institutions is rapidly growing.

FAO, which was the steward of VB2, migrated to VB3 for the maintenance of their thesaurus Agrovoc in August 2018 and has recently started adopting it for classification systems used in statistics. INRA, the French "Institut national de la recherche agronomique" [26] and CIRAD, "la recherche agronomique pour le développement"[27] showed their intention to move to VB3 (from VB2 and as a first adoption in respectively) for the management of their thesauri. Still in the field of agriculture, the US National Agricultural Library (NAL) [28], together with FAO and CABI ("Centre for Agriculture and Bioscience International")[29] agreed to use the VB3 platform in the context of the GACS[30] project, a global agriculture thesaurus born from the integration of the three respective thesauri (NALT, Agrovoc and the CAB thesaurus) of these major players in the area.

The Senate of the Italian Republic is also migrating management of its thesaurus Teseo from VB2 to VB3. There are also newcomers who are already adopting the platform and started directly with version 3, such as GelbeSeiten[31], the German Yellow Pages, which are


Figure 14. Download statistics for VocBench 2 (on the left) and VocBench 3 (on the right) as of 8 February 2019

using VB3 to maintain their homonymous thesaurus, and the Solidaridad Network[32].

Furthermore, VB3 acquired much more visibility and potential for adoption with respect to its predecessors, since it became the reference platform of the EU[33], recommended by the European Commission to the member states for the collaborative management of thesauri, ontologies and (now) lexicons.

To improve our understanding of the uptake of VocBench, we collected some statistics from the download pages of VB2[34] and VB3[35] as of 8th of February 2019. Figure 14 draws how many times each release of VB2 (on the left), respectively, of VB3 (on the right) was downloaded.

The most downaloded version of VB2 is the 2.3 with 395 downloads. The lower statistics associated with subsequent releases can be explained by the fact they were released one after the other with only a few months in between. These more frequent releases allowed us to deliver new features and (important) fixes to users; however, not every user had the same incentives to keep update with this faster release candence. The last release of VB2 2.4.4 (available since 8 March 2017) registered 295 downloads.

On the download page of VB2, there are also the documentations of VB 2.3 and VB 2.1 in PDF format, which have in total 959 downloads. A related obseravation is that the two versions of a simple thesaurus used in the getting started documentation were downloaded 491 times. Perhaps more important is the fact that a patch for a problem with the validation in VB 2.3 was downloaded 149 times: this figures tells us about users who considered the information they

were validating valuable enough to warrant the application of the patch.

On the right side of Figure 14, we report analogous statistics for VB3. The first stable release (version 1.0) has 138 downloads, while the subsequent version 2.0 was downloaded 498 times (more than the most downloaded version of VB2). Version 3.0.1 (released five days after 3.0.0) has 438 downloads, while version 4.0.2 (released three days after 4.0.0) has 637 downloads being the most downloaded so far. These high figures are associated with relatively frequent releases (every few months), and thus support two (non-mutually exclusive) hypotheses: existing users are willing to test/use new versions as they bring substantial new features, or there is a sufficiently large supply of new users. The typical slowness of large organizations (which are typical adopters of VocBench) in reacting to changes (confirmed by the feedback we have gathered, where most of them declared to have "intention to move" or "being in the process of migrating" though still with an unclear schedule) suggests that the second hypothesis covers at least a non-trivial part of the phenomenon.

The user community's main point of interaction and support is the VocBench discussion group[36]. The group (as of 8th February 2019) counts 154 members. In order to roughly quantify the impact of the group, we compared its size with the analogous community forum of RDF4J[37]. The user group of RDF4J counts 390 members, so the sizes of the two groups are in the same order of magnitude. We must also consider that RDF4J is one of the two most popular Java RDF

---

[32] https://www.solidaridadnetwork.org/

[33] https://ec.europa.eu/isa2/solutions/vocbench3_en

[34] https://bitbucket.org/art-uniroma2/vocbench2/downloads/

[35] https://bitbucket.org/art-uniroma2/vocbench3/downloads/

[36] http://groups.google.com/group/vocbench-user

[37] http://groups.google.com/group/rdf4j-users

frameworks (the other is Jena[38]), thus relying on a less scattered user distribution than RDF editors may have.

User questionnaires are occasionally sent to the mailing list of VocBench in order to gather feedback from users about satisfaction and ideas for new directions for the system. A recent investigation held in March, 2018, based on a questionnaire diffused on the mailing list revealed that roughly 50% of the respondents already adopted VocBench 3 in production, 22% were in the process of adopting it, 17% were already using VocBench 2 in production and were considering moving to VocBench 3 while 11% was shared by users adopting VB2 and still not sure whether to move to the new platform and newcomers evaluating the VB3 platform. Adoption of VocBench 3 is still mainly focused on thesaurus development: among those adopting VB3 in production, 89% of the respondents are using it for developing thesauri while 33% is adopting it for OWL ontology development. We do not have figures for OntoLex development as at the time of the investigation the OntoLex editor in VB had not yet been released.

6. Related Work

VB3 has dedicated facilities for a variety of core RDF knowledge models, and it can be used both as a desktop tool and as a web-based collaborative environment. The number of systems it could be compared against is consequently very large. Instead of attempting an exhaustive survey of these systems, we focused on selected systems, which are currently and widely recognized as references in their area of specialization.

6.1. Protégé

Protégé[39] [43] is an open source, very popular ontology editor. Our use of the term Protégé avoids ambiguity with the web-based version WebProtégé [44].

Originally based on a *frame* language (with a later extension for OWL), Protégé 4 switched to OWL-DL through the OWL API [45]. Protégé can thus benefit from DL reasoners, such as HermiT [46] which combines completeness, feature-richness and scalability.

DL semantics precluded most meta-modelling, which is supported by RDF systems (such as VB3), based on OWL Full semantics and rule-based reasoning. OWL2 closed this gap with the introduction of

punning. Regarding the relative strengths of each approach, tableau-based (DL) reasoners are usually less scalable and efficient than lightweight rule-based reasoners, which, however, are less expressive [47].

One of the strengths of Protégé lies in its extensibility through plugins. At the time of writing, its plugin library[40] contains 120 items. One of them is SKOSed [48], consisting in "a suite of views and tools for working specifically with SKOS". SKOSed adds convenience (e.g. the hierarchical representation of concepts) on top of the foundation laid down by Protégé, which guarantees the possibility to intermix SKOS and OWL (e.g. to extend SKOS with additional classes or properties), and to leverage already existing features (e.g. reasoners). Under this viewpoint, SKOSed is very close to how VB3 achieves its multi-model support, even though SKOSed offers a notably smaller set of functionalities with respect to VocBench and is limited to the desktop version of Protégé. Moreover, SKOSed is no longer maintained, and the last version is declared compatible with Protégé 4.1+: a quick test revealed that SKOSed is incompatible with Protégé 5.2.

Beyond single-user usage, Protégé supports a client-server mode (or multiuser model). Different instances of Protégé (clients) connect to a shared instance (server), enabling users to co-work on the same ontology (or SKOS dataset, in case of SKOSed). Client-server communication is based on Java RMI, while VB3 uses HTTP and JSON (popular in Web APIs).

The client-server mode greatly differs between Protégé 3 and Protégé 4: in Protégé 3 individual changes are sent to the server as they are made by clients, while in Protégé 4 changes are grouped together and only sent when the user commits them[41]. In Protégé 3, the clients cannot work without a connection to the shared server. This scenario is conceptually similar to a web application (in particular a single-page application, such as VB3), although clients are locally installed rather than downloaded on the fly from the web. Conversely, in Protégé 4, the clients can work offline, as the connection to the server is only required by server-related operations, such as commit or update. In fact, the ontology can even be modified outside of Protégé: upon commit, Protégé will compute the difference between the current version and the one originally fetched from the server. The client-server mode of Protégé 4 is thus conceptually close to a centralized version control system specialized for ontologies.

---

[38] http://jena.apache.org/
[39] http://protege.stanford.edu/
[40] https://protegewiki.stanford.edu/wiki/Protege_Plugin_Library

[41] https://protegewiki.stanford.edu/wiki/Protege4ClientServer

Collaborative Protégé [49] is an extension for Protégé 3 that introduces several collaboration-oriented features beyond the multiuser model (client-server mode). Noteworthy features include the annotation of ontology components and changes, the management of a change history, inline discussions and a chat. In VB3, history is managed internally, while other capabilities are delegated to external (already appreciated) services (e.g. JIRA for issue management), through extensible connectors, while offering a smooth integration (e.g. showing issues for a resource inside its resource-view).

6.2. VoCol

Version control systems for software development (e.g. GIT) are often considered inadequate for versioning of ontologies and RDF datasets: operating on a serialization of an ontology, they can be confused by semantically irrelevant operations (e.g. statement reordering, change of blank node identifiers).

Nonetheless, VoCol [50,51] demonstrated the viability of GIT[42] for distributed development of vocabularies, by proposing a methodology and a suite of integrated tools. VoCol orchestrates the execution of different tasks (e.g. validation, documentation generation, etc.) in response to notifications sent from a GIT repository, often in response to events such as pushing of new commits. VoCol is not an ontology editor on its own. Actually, the wiki of VoCol suggests editing the files versioned with GIT using a text editor, since some ontology/RDF editors tend to completely change the serialization of the data (e.g. the order of the triples) upon saving the updated version[43].

6.3. WebProtégé

WebProtégé was initially thought as a lightweight (open source) ontology editor, but it has grown over time reducing the functional gap with Protégé. Indeed, it supports class expressions in the Manchester syntax (as VB3) and editing of SWRL rules (absent in VB3).

WebProtégé reused and extended the functionalities of Collaborative Protégé for change tracking, notes, discussions and access control. In WebProtégé, users can also monitor ("watch") individual resources or entire branches of the class tree: changes to these resources are listed in the user interface or sent via email notifications. This feature – not yet supported by VB3 – may speed up interactions between editors co-working on the same area of the ontology.

In addition to the change history of the ontology, WebProtégé enables to download the ontology in each state of its evolution. This is a rather unique feature, since other systems (VB3 included) only supports discrete snapshots of the edited data. A similar behavior may be found in VoCol, since in GIT it is possible to download any (previous) state of a repository.

WebProtégé supports the acquisition of instance data by means of customizable forms that effectively shield domain experts from the complexity of the underlying ontological modelling. This mechanism was borrowed from Protégé 3, and it is somewhat related to VB3 custom forms. Actually, VB3 custom forms are meant to enhance instance construction or property setting, without the aim to replace the editing panel (i.e. the resource-view) as a whole.

Currently, WebProtégé does not support plugins like its desktop counterpart[44]; however, its modular architecture makes it easy to modify the code of the system and introduce additional components (e.g. custom tabs or custom form widgets). Conversely, VB3 already supports plugins at the service level, with the dynamic generation of configuration dialogs in the user interface. Nonetheless, provisioning of arbitrary extensions for the user interface (as available in Protégé desktop) is still missing in VB3, due to limitations of the Angular technology.

In line with its use in the life science domain, WebProtégé provides a dedicated facility for linking terms from BioPortal[45] [52] ontologies. Starting from version 5.0, VB3 is going support data catalogs in general, and provide interesting features based on them (e.g. importing an ontology from a catalog).

Currently, WebProtégé lacks dedicated facilities for SKOS, nor can third-party plugins like SKOSed fill this gap. Fortunately, there are numerous alternative SKOS editors, most of which – especially the most recent ones – are web-based collaborative editors. Mochón *et al.* surveyed [53] existing thesaurus managements tools from the perspective of LOD. They found that PoolParty[46] outperforms other systems with respect to every indicator, while TemaTres[47] [54] ranked second. In the following paragraphs, we will discuss these two systems, as well a few others (also included

---

[42] https://git-scm.com/

[43] https://github.com/vocol/vocol/wiki/Developing-Vocabularies-with-VoCol-Environment

[44] https://mailman.stanford.edu/pipermail/protege-user/2018-March/008741.html

[45] https://bioportal.bioontology.org/

[46] https://www.poolparty.biz/

[47] https://www.vocabularyserver.com/

in the survey), which we considered noteworthy either for impact or for diversity of approach.

6.4. PoolParty

PoolParty is a highly regarded suite of semantic technologies for different needs, including i) thesaurus management, ii) term and phrase extraction, entity linking and disambiguation, and iii) semantic search and data mining. Different combinations of these proprietary technologies are available in different combinations as a one-time purchase or as a renewable subscription. We evaluated a demo instance of PoolParty Enterprise Server, which combines thesaurus management and term extraction, linking and disambiguation. The costlier PoolParty Semantic Integrator also includes semantic search & data mining capabilities and a wider spectrum of triple store backends. The latter is supported by VB3 by a dedicated extension point.

In PoolParty, the analogous of the resource-view for a concept consists of several tabs corresponding to various functionalities, such as detailed editing, notes, history, quality management, etc. The editing tab itself is subdivided into tabs, corresponding to different cohesive groups of properties. By default, it is possible to use the SKOS tab to edit the core description of a concept. However, additional tabs can be activated, by choosing an ontology or a custom scheme (a view on a subset of one or more ontologies developed for some purpose). Indeed, PoolParty treats ontologies as second-class citizens: while it has long been possible to edit them inside PoolParty, until version 7, editing of ontologies used a completely different, very constrained user interface. Indeed, only PoolParty 7 recently introduced a class tree for the class taxonomy.

PoolParty supports change tracking, while history can be browsed both globally and per resource. The approval workflow allows to exert some control on proposed changes: concepts transition to the draft state upon any change, and an explicit action is required to transition them again to the approved state. Upon unapproved changes, the concept can be assigned to a user for remedy. If the thesaurus is seriously compromised, it is possible to revert to a previously snapshotted state. VB3 also supports snapshots and quickly switching to a snapshot for visualization. However, reverting to a snapshotted state is – at the time of writing – not supported. PoolParty supports quality management through the integration of qSKOS[48] [55], which

can evaluate several quality criteria (which inspired our constraint validation subsystem). These criteria can be enforced interactively (i.e. forbid a request violating a constraint) or only included in quality reports.

PoolParty can leverage existing Linked Datasets, e.g. to populate a thesaurus or create mappings (e.g. to DBpedia). Similarly, VB3 supports browsing (via the resource-view) remote resources (accessed via HTTP or SPARQL) and creating mappings to remote datasets.

PoolParty supports (as VB3 does) the association of JIRA for task management.

Beyond thesaurus management in strict sense, interesting features of PoolParty (in some offerings) are the acquisition of concepts/terms from a corpus and the semantic annotation/classification of documents. Currently, VB3 lacks these capabilities. One may argue that they should be provided by external services that (right now) can interact with VB3 through its API or via a shared triple store. Indeed, there have already been some attempts to integrate knowledge acquisition from text in our platform [56,57] through CODA.

6.5. TopBraid EVN

TopBraid EVN[49] (Enterprise Vocabulary Net) is a proprietary web-based editor for thesauri, ontologies, corpora, content tag sets and crosswalks. Differently from PoolParty, TopBraid EVN supports ontologies to the same extent as thesauri. Moreover, the management of these diverse assets is in fact very uniform, at least for what concerns metadata management, access policy definition, import/export, etc. The main differences arise in the editing section, where unsurprisingly ontologies and thesauri use a different interface than, for example, crosswalks. Unless stated otherwise, we will refer to thesaurus/ontology editing.

TopBraid EVN implements change governance through working copies. Alike branches in GIT, working copies can be edited autonomously, and then merged to the production version of the asset, if the proposed changes are approved. The approval process is supported by a comparison report showing the difference between the copy and the production version in terms of property value changes of resources. The same type of report is available to editors, to understand how a resource has evolved, and in case undo a change. In addition to this resource-centric view, TopBraid EVN supports browsing and searching the history in terms of coarse-grained changes consisting of

---

multiples triple additions/removals. Again, it is possible to revert individual changes. However, reverting a change in the middle of the history may create orphan resources. Validation in VB3 has the same problem. The behavior of TopBraid EVN differs from the validation in VB3 with respect to their interaction with the history. Reverting a change appends to the history another change that undoes the effects of the former. In VB3, a change pending for validation is – on purpose – not logged in the history, and if it is rejected, it is like the change had never happened. In TopBraid EVN, working copies enable a similar result. Archiving a working copy is alike snapshots in PoolParty and VB3.

TopBraid EVN implements integrity checks via SPIN[50] (SPARQL Inference Notation) and, more recently, via SHACL (Shapes Constraint Language) [58]. These constraints can be evaluated interactively or included in a report (like PoolParty). The use of a standard, declarative language eases adding constraints. In VB3, constraints are hardcoded with some duplication between the checks within services and the integrity constraint validation subsystem: this approach enables dedicated visualizations and different, contextualized quick fixes to ease debugging and repair of anomalies.

Like WebProtégé, TopBraid EVN allows to customize the form associated with a class. Internally represented in SWP[51] (SPARQL Web Pages), these forms can be edited (to an extent) with a graphical editor.

TopBraid EVN supports the assignment of tasks to users at different levels (e.g. asset or individual resources). Additionally, TopBraid EVN supports comments on resources. Transitioning between a few statuses, comments also provide a simple task system.

In addition to these integrated facilities for collaboration, TopBraid EVN supports linking projects to JIRA: e.g. the editing view focused on a resource enables to see or create related issues on JIRA. As already stated, PoolParty and VB3 also offer this capability. Actually, VB3 defines an extension point that can be implemented for different collaboration platforms.

Like PoolParty, TopBraid EVN can manage document corpora, and supports both manual and automatic (document-level) semantic tagging of documents.

6.6. TemaTres

TemaTres is an open source web application for the management of thesauri, glossaries, controlled vocab-

ularies, etc. TemaTres uses a term-based model in contrast to the concept-based model underpinning SKOS. Indeed, SKOS is only available as an import/export format, while data is stored in a relational DB using a dedicated schema. This schema is enough flexible to support custom relations; however, it doesn't support the import of custom ontologies, and these customizations do not make it into the SKOS export anyway.

TemaTres supports multilingualism via interlinked monolingual vocabularies, either on the same TemaTres instance or accessed via a terminological web service, enabling the federation of TemaTres instances. Federation also benefits vocabulary mapping, allowing for term-lookup over remote vocabularies. VB3 achieves a similar goal through "assisted search", enabling term lookup over the SPARQL endpoints of open online datasets. TemaTres supports bulk generation of mappings, by fetching and then reversing mappings contained in a federated vocabulary.

TemaTres has a workflow model like the one of VB3: concepts are created in a candidate state, and once accepted, they do not reenter the "modified" state after each modification. While TemaTres implements the state as an explicit metadata of the term (as in VB2), in VB3 the state is represented implicitly through the state of the triple asserting the characterizing class of the resource (when validation is enabled).

TemaTres has a very limited support for change history in the form of recent changes to a vocabulary (also available as an RSS feed).

TemaTres can be configured to enforce some integrity constraints (e.g. disallow polyhierarchy), while offering some quality assurance tools (e.g. terms without hierarchical relationships). The TemaTres Keyword Distiller uses terms extracted from text[52] to obtain keywords from a controlled vocabulary.

6.7. iQvoc

iQvoc[53] is a vocabulary management system based on SKOS originated from a research on the challenges and opportunities of Linked Data publication via web frameworks. iQvoc uses the framework Ruby on Rails. Despite an initial binding to a specific vocabulary, subsequent development focused on separating the core logic from vocabulary-specific customizations. When iQvoc 3.0 was open sourced, even the support

---

for SKOS-XL was moved to a separate module[54]. iQvoc is like TemaTres under several viewpoints, such as resources being displayed as directly addressable content-pages, federation, bulk-generation of mappings, etc. However, iQvoc differs profoundly from TemaTres in being natively based on SKOS. iQvoc does not work at the level of RDF statements, but SKOS is implemented as a domain model persisted to a relational DB via the ORM (Object-Relational Mapping) found in Ruby on Rails. Extensions to the domain model must be realized as new code modules.

Governance of change in iQvoc is based on the idea of differentiating the published version of a concept and its (unique) draft version created upon a change to that concept. Furthermore, locking the draft version prevents conflicts. Users with suitable rights can publish the draft version of a concept, making the changes persistent and visible. Before publishing, these users can request the evaluation of some integrity constraints, in order to avoid the corruption of the data being edited. In iQvoc, the granularity of validation coincides with individual concepts being edited, while VB3 records individual operations for validation. The approach of iQvoc enables to group together individual changes that are best vetted together. However, it is impossible to validate cross-concept changes.

iQvoc lacks a global history of changes; however, changes to individual resources can be described by attaching SKOS change notes to the affected resource (creator and date time are set automatically).

6.8. Lemon Editors

To our knowledge, VB3 is the first multi-model editor concerned with OntoLex-Lemon. In literature we could only find purpose-built systems that are tailored to a specific application of *lemon* (in broad sense, thus including the legacy Monnet *lemon* [9]). McCrae *et al* argued [59] that generic data-driven editors are difficult to use with the *lemon* model: some model constructs are not properly displayed, and, moreover, there are constructs made of a few triples that logically should be manipulated as a single entity. McCrae *et al* addressed these issues through the development of Lemon Source. Lemon Source also addresses collaboration, which was identified as a key requirement for lexicon development. The system thus incorporates change tracking, role-based access control, validation

workflow, etc. Most of these items are generally useful features, already found in general-purpose collaborative editors such as VB3. Both Lemon Source and VB3 hide (most of) the complexity of the model, by prompting the user for essential information about a given construct (say the lemma of a lexical entry), which drives the generation of its complex RDF serialization. Lemon Source automates the construction of ontology lexicons a bit further: by means of a configurable NLP (Natural Language Processing) pipeline, it creates lexical entries from (RDFS) labels commonly found in ontologies, computing their tokenization, syntactic tree, and using (whenever possible) entries in already existing lexicons (e.g. Princeton WordNet). Unfortunately, Lemon Source is no longer publicly available.

Lemon Assistant [60] is another web-based editor for *lemon*, which is based on the design patterns for ontology lexicons [21]. Its motivation lies in the observation that the use of *lemon* is complex and error prone. In VB3, we implemented all the patterns supported by Lemon Assistant, but we complemented them with more general operations on *lemon* models, so that we are not constrained to ontology lexicons. Indeed, despite what the extended name ("lexicon model for ontologies") of the model would suggest, *lemon* is largely used beyond its original scope, for example for the representation of language resources as Linked Data. However, VB3 does not support – at the moment – the generation of usages examples of lexical entries and some integrity (cross-language) checks. Lemon Assistant is currently available only as an online service[55], no longer integrated with the integrity checks.

LexO [61,62] is the most recent *lemon* editor, grown in the area of Digital Humanities to support philologists, historical linguists and lexicographers. There are different versions of LexO that were developed in different projects to support different linguistic phenomena through a combination of extensions of the *lemon* model and customizations of the user interface: the notion of root for Arabic, transliteration and tone for Chinese, multiple alphabets, multi word expressions combining words in different languages and ancients forms without a canonical form (because were never attested) for Old-Occitan. Conversely, the demo[56] linked by the project page[57] shows only a very basic support for ontologies (just the class tree, without the possibility to edit the details of the classes).

In Monnet *lemon*, the representation of a wordnet was achieved as if it were a lexicalization of an ontology: synsets were usually represented as instances of an arbitrary class, while words were represented as lexical entries. Previous editors supported the editing of language resources with this mechanism. However, OntoLex-Lemon eventually differentiated between an ontology lexicon and a wordnet-like resource: the class *lexical concept* models synsets, while the attachment of lexical entries is achieved with a dedicated set of properties. Obviously, editors developed long before OntoLex-Lemon may not support all of this. Conversely, VB3 features an extensive support for lexical concepts, and it was tested with very large resources: for example, in Section 4.1.4 we cited the possibility to manage the collection of 34 wordnets collected in Open Multilingual Wordnet. Another recent innovation of the model not supported by previous editors is the LIME metadata module. VB3 supports the computation and the export of this kind of metadata. In the future, VB3 will benefit from this metadata to better orchestrate the alignment of different datasets: e.g. determine the overlap in the supported natural languages, or the identification of useful language resources.

7. Conclusion and Future Work

In the last years, VocBench has addressed the needs of large organizations, companies and independent users needing an open source collaborative environment for editing thesauri, supporting a formalized editorial workflow. Continuous user feedback allowed us to spot bugs and to improve the usability of VocBench.

It is thanks to this community feedback, to the support of the ISA^2 program and to our desire to reach new quality levels that we started this endeavor, by rethinking most of VocBench from scratch, still benefiting from the experiences we had with VB2.

The most important achievement of the new platform lies at its core: a fully-fledged RDF core framework, developed by further improving the Semantic Turkey framework with functionalities for user management, role-based access control, change tracking and collaboration and by providing it with a new user interface. With respect to VB2, we could say VB3 is a user interface (delivered as a web application) for the improved Semantic Turkey platform. A second major achievement is the broadened support for all major standards of the RDF family, going beyond SKOS thesauri and embracing OWL ontologies and OntoLex (both as a core model for developing lexicons and as a

lexicalization model for all kind of RDF datasets) which makes of VB a unique, comprehensive offer in the scenario of RDF development platforms.

We hope that this evolution of the system will lay a solid foundation for the realization of a new range of services spacing from knowledge acquisition, evolution and management in Europe and worldwide.

Acknowledgments

References

[1] C. Caracciolo, A. Stellato, A. Morshed, G. Johannsen, S. Rajbhandari, Y. Jaques, and J. Keizer, "The AGROVOC Linked Dataset," *Semantic Web Journal*, vol. 4, no. 3, pp. 341–348, 2013, doi: 10.3233/SW-130106 .

[2] A. Stellato, S. Rajbhandari, A. Turbati, M. Fiorelli, C. Caracciolo, T. Lorenzetti, J. Keizer, and M.T. Pazienza, "VocBench: a Web Application for Collaborative Development of Multilingual Thesauri," in *The Semantic Web. Latest Advances and New Domains (Lecture Notes in Computer Science)*.: Springer International Publishing, 2015, vol. 9088, pp. 38-53, doi: 10.1007/978-3-319-18818-8_3 .

[3] World Wide Web Consortium (W3C). (2009, August) World Wide Web Consortium (W3C). [Online]. Available: http://www.w3.org/TR/skos-reference/

[4] World Wide Web Consortium (W3C). (2009, August) World Wide Web Consortium (W3C). [Online]. Available: http://www.w3.org/TR/skos-reference/skos-xl.html

[5] G. Hodge, *Systems of Knowledge Organization for Digital Libraries: Beyond Traditional Authority Files*. Washington, DC: Council on Library and Information Resources, April 2000. [Online]. Available: https://www.clir.org/pubs/reports/pub91/

[6] M. T. Pazienza, N. Scarpato, A. Stellato, and A. Turbati, "Semantic Turkey: A Browser-Integrated Environment for Knowledge Acquisition and Management," *Semantic Web Journal*, vol. 3, no. 3, pp. 279-292, 2012, doi: 10.3233/SW-2011-0033 .

[7] D. Griesi, M. T. Pazienza, and A. Stellato, "Semantic Turkey - a Semantic Bookmarking tool (System Description)," in *4th European Semantic Web Conference (ESWC 2007)*, Innsbruck, Austria, 2007, doi: 10.1007/978-3-540-72667-8_56 , June 3-7.

[8] A. Stellato, A. Turbati, M. Fiorelli, T. Lorenzetti, E. Costetchi, C. Laaboudi, W. Van Gemert, and J. Keizer, "Towards VocBench 3: Pushing Collaborative Development of Thesauri and Ontologies Further Beyond," in *17th European Networked Knowledge Organization Systems (NKOS) Workshop. Thessaloniki, Greece,*

September 21st, 2017, 2017, pp. 39-52. [Online]. Available: http://ceur-ws.org/Vol-1937/paper4.pdf

[9] J. McCrae, D. Spohr, and P. Cimiano, "Linking Lexical Resources and Ontologies on the Semantic Web with Lemon," in *The Semantic Web: Research and Applications (Lecture Notes in Computer Science)*.: Springer Berlin Heidelberg, 2011, vol. 6643, pp. 245-259, doi: 10.1007/978-3-642-21034-1_17 .

[10] J. P. McCrae, J. Bosque-Gil, J. Gracia, P. Buitelaar, and P. Cimiano, "The OntoLex-Lemon Model: Development and Applications," in *Electronic lexicography in the 21st century. Proceedings of eLex 2017 conference.*, 2017, pp. 587-597. [Online]. Available: https://elex.link/elex2017/wp-content/uploads/2017/09/paper36.pdf

[11] M. Fiorelli, A. Stellato, T. Lorenzetti, A. Turbati, P. Schmitz, E. Francesconi, N. Hajlaoui, and B. Batouche, "Towards OntoLex-Lemon editing in VocBench 3," *AIDAinformazioni*, vol. 36, no. special issue, pp. 81-102, 2018, doi: 10.4399/97888255216345 .

[12] M. Fiorelli, A. Stellato, T. Lorenzetti, A. Turbati, P. Schmitz, E. Francesconi, N. Hajlaoui, and B. Batouche, "Editing OntoLex-Lemon in VocBench 3," in *Proceedings of the Twelfth International Conference on Language Resources and Evaluation (LREC 2020)*, 2020, (in press).

[13] P. Jain, P. Hitzler, P. Z. Yeh, K. Verma, and A. P. Sheth, "Linked Data Is Merely More Data," in *Linked Data Meets Artificial*. Menlo Park: AAAI Press, 2010, pp. 82–86. [Online]. Available: https://www.aaai.org/ocs/index.php/SSS/SSS10/paper/view/1130/1454

[14] B. Bishop, A. Kiryakov, D. Ognyanoff, I. Peikov, Z. Tashev, and R. Velkov, "OWLIM: A family of scalable semantic repositories," *Semantic Web*, vol. 2, no. 1, pp. 33-42, 2011, doi: 10.3233/SW-2011-0026 .

[15] M. Fiorelli, M. T. Pazienza, A. Stellato, and A. Turbati, "CODA: Computer-aided ontology development architecture," *IBM Journal of Research and Development*, vol. 58, no. 2/3, pp. 14:1 - 14:12, March-May 2014, doi: 10.1147/JRD.2014.2307518 .

[16] M.T. Pazienza, A. Stellato, and A. Turbati, "PEARL: ProjEction of Annotations Rule Language, a Language for Projecting (UIMA) Annotations over RDF Knowledge Bases," in *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC'12)*, Istanbul, Turkey, 2012. [Online]. Available: http://www.lrec-conf.org/proceedings/lrec2012/pdf/424_Paper.pdf

[17] M. Fiorelli, T. Lorenzetti, M. T. Pazienza, A. Stellato, and A. Turbati, "Sheet2RDF: a Flexible and Dynamic Spreadsheet Import&Lifting Framework for RDF," in *Current Approaches in Applied Artificial Intelligence*, M. Ali et al., Eds.: Springer International Publishing, 2015, vol. 9101, pp. 131-140, doi: 10.1007/978-3-319-19066-2_13 .

[18] F. Cotton, R. Cyganiak, R.T.A.M. Grim, D. W. Gillman, Y. Jaques, and W. Thomas, "XKOS: An SKOS Extension for Statistical Classifications," in *Proceedings 59th ISI World Statistics Congress, 25-30 August 2013, Hong Kong (Session CPS203)*, 2013, pp. 5492-5497. [Online]. Available: http://2013.isiproceedings.org/Files/CPS203-P32-S.pdf

[19] C. Fellbaum, *WordNet: An Electronic Lexical Database*. Cambridge, MA: WordNet Pointers, MIT Press, 1998. [Online]. Available: https://mitpress.mit.edu/books/wordnet

[20] M. Fiorelli, T. Lorenzetti, M. T. Pazienza, and A. Stellato, "Assessing VocBench Custom Forms in Supporting Editing of Lemon Datasets," in *Language, Data, and Knowledge (Lecture Notes in Artificial Intelligence)*.: Springer, Cham, 2017, vol. 10318, pp. 237-252, doi: 10.1007/978-3-319-59888-8_21 .

[21] J. P. McCrae and C. Unger, "Design Patterns for Engineering the Ontology-Lexicon Interface," in *Towards the Multilingual Semantic Web*, P. Buitelaar and P. Cimiano, Eds.: Springer Berlin Heidelberg, 2014, pp. 15-30, doi: 10.1007/978-3-662-43585-4_2 .

[22] M. Heller. (2007, January) InfoWorld. [Online]. Available: https://www.infoworld.com/article/2640739/application-development/rest-and-crud--the-impedance-mismatch.html

[23] I. Bratko, *Prolog Programming for Artificial Intelligence*.: Addison Wesley, 2001. [Online]. Available: https://books.google.it/books?id=-15su78YRj8C

[24] E. Denti, A. Omicini, and A. Ricci, "tuProlog: A Light-Weight Prolog for Internet Applications and Infrastructures," in *Practical Aspects of Declarative Languages (Lecture Notes in Computer Science)*.: Springer, Berlin, Heidelberg, 2001, vol. 1990, pp. 184-198, doi: 10.1007/3-540-45241-9_13 .

[25] M. Fiorelli, M. T. Pazienza, A. Stellato, and A. Turbati, "Version Control and Change Validation for RDF Datasets," in *Metadata and Semantic Research (Communications in Computer and Information Science)*, E. Garoufallou et al., Eds.: Springer, Cham, 2017, vol. 755, pp. 3-14, doi: 10.1007/978-3-319-70863-8_1 .

[26] M. Fiorelli, M. T. Pazienza, and A. Stellato, "Change management and validation for collaborative editing of RDF datasets," *International Journal of Metadata, Semantics and Ontologies*, vol. 12, no. 2-3, 2017, doi: 10.1504/IJMSO.2017.090783 .

[27] D. Hernández, A. Hogan, and M. Krötzsch, "Reifying RDF: What Works Well With Wikidata?," in *Proceedings of the 11th International Workshop on Scalable Semantic Web Knowledge Base Systems, Bethlehem, PA, USA, October 11, 2015*, 2015, pp. 32-47. [Online]. Available: http://ceur-ws.org/Vol-1457/SSWS2015_paper3.pdf

[28] V. Nguyen, O. Bodenreider, and A. Sheth, "Don't Like RDF Reification?: Making Statements About Statements Using Singleton Property," in *Proceedings of the 23rd International Conference on World Wide Web, April 7-11, 2014, Seoul, South Korea*, 2014, pp. 759-770, doi: 10.1145/2566486.2567973 .

[29] J. J. Carroll, C. Bizer, P. Hayes, and P. Stickler, "Named Graphs, Provenance and Trust," in *WWW '05: Proceedings of the 14th international conference on World Wide Web*, New York, NY, USA, 2005, pp. 613-622, doi: 10.1145/1060745.1060835 .

[30] R. Laurens and H. Rinkea, "The YASGUI family of SPARQL clients," *Semantic Web*, vol. 8, no. 3, pp. 373-383, 2017, doi: 10.3233/SW-150197 .

[31] C. Bizer, J. Lehmann, G. Kobilarov, S. Auer, C. Becker, R. Cyganiak, and S. Hellmann, "DBpedia - A crystallization point for the Web of Data," *Web Semantics: Science, Services and Agents on the World Wide Web*, vol. 7, no. 3,

pp. 154–165, September 2009, doi: 10.1016/j.websem.2009.07.002 .

[32] J. David, J. Euzenat, F. Scharffe, and C. Trojahn dos Santos, "The Alignment API 4.0," *Semantic Web Journal*, vol. 2, no. 1, pp. 3-10, 2011, doi: 10.3233/SW-2011-0028 .

[33] World Wide Web Consortium (W3C). (2014, January) World Wide Web Consortium (W3C). [Online]. Available: http://www.w3.org/TR/vocab-dcat/

[34] M. Dekkers. (2013, August) World Wide Web Consortium (W3C). [Online]. Available: http://www.w3.org/TR/vocab-adms/

[35] K. Alexander, R. Cyganiak, M. Hausenblas, and J. Zhao. (2011, March) World Wide Web Consortium (W3C). [Online]. Available: http://www.w3.org/TR/void/

[36] M. Fiorelli, A. Stellato, J. P. Mccrae, P. Cimiano, and M. T. Pazienza, "LIME: the Metadata Module for OntoLex," in *The Semantic Web. Latest Advances and New Domains (Lecture Notes in Computer Science)*.: Springer International Publishing, 2015, vol. 9088, pp. 321-336, doi: 10.1007/978-3-319-18818-8_20 .

[37] M. Fiorelli, M. T. Pazienza, and A. Stellato, "An API for OntoLex LIME datasets," in *OntoLex-2017 1st Workshop on the OntoLex Model (co-located with LDK-2017)*, Galway, 2017. [Online]. Available: http://ceur-ws.org/Vol-1899/OntoLex_2017_paper_8.pdf

[38] M.T. Pazienza, S. Sguera, and A. Stellato, "Let's talk about our "being": A linguistic-based ontology framework for coordinating agents," *Applied Ontology, special issue on Formal Ontologies for Communicating Agents*, vol. 2, no. 3-4, pp. 305-332, December 2007. [Online]. Available: https://content.iospress.com/articles/applied-ontology/ao035

[39] M.T. Pazienza and A. Stellato, "An Environment for Semi-automatic Annotation of Ontological Knowledge with Linguistic Content," in *The Semantic Web: Research and Applications (Lecture Notes in Computer Science)*.: Springer, 2006, vol. 4011, pp. 442-456, doi: 10.1007/11762256_33 .

[40] M. Fiorelli, M. T. Pazienza, and A. Stellato, "A Meta-data Driven Platform for Semi-automatic Configuration of Ontology Mediators," in *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, Reykjavik, Iceland, May 2014. [Online]. Available: http://www.lrec-conf.org/proceedings/lrec2014/pdf/1059_Paper.pdf

[41] M. Fiorelli, A. Stellato, T. Lorenzetti, P. Schmitz, E. Francesconi, N. Hajlaoui, and B. Batouche, "Metadata-driven Semantic Coordination," in *Metadata and Semantic Research (Communications in Computer and Information Science)*, E. Garoufallou, F. Fallucchi, and E. William De Luca, Eds.: Springer, Cham, 2019, vol. 1057, doi: 10.1007/978-3-030-36599-8_2 .

[42] E. Francesconi, M.W. Küster, P. Gratz, and S. Thelen, "The Ontology-based Approach of the Publications Office of the EU for Document Accessibility and Open Data Services," in *International Conference on Electronic Government and the Information Systems Perspective (EGOVIS 2015)*, Valencia, Spain, 2015, doi: 10.1007/978-3-319-22389-6_3 .

[43] M. A. Musen, "The Protégé Project: A Look Back and a Look Forward," *AI Matters*, vol. 1, no. 4, pp. 4-12, 2015, doi: 10.1145/2757001.2757003 .

[44] T. Tudorache, C. Nyulas, N. F. Noy, and M. A. Musen, "WebProtégé: A Collaborative Ontology Editor and Knowledge Acquisition Tool for the Web," *Semantic Web*, vol. 4, no. 1, pp. 89-99, 2013, doi: 10.3233/SW-2012-0057 .

[45] M. Horridge and S. Bechhofer, "The OWL API: A Java API for OWL ontologies," *Semantic Web*, vol. 2, no. 1, pp. 11-21, 2011, doi: 10.3233/SW-2011-0025 .

[46] B. Glimm, I. Horrocks, B. Motik, G. Stoilos, and Z. Wang, "HermiT: An OWL 2 Reasoner," *Journal of Automated Reasoning*, vol. 53, no. 3, pp. 245-269, 2014, doi: 10.1007/s10817-014-9305-1 .

[47] A. Hogan, J. Z. Pan, A. Polleres, and Y. Ren, "Scalable OWL 2 Reasoning for Linked Data," in *Reasoning Web. Semantic Technologies for the Web of Data. Reasoning Web 2011 (Lecture Notes in Computer Science)*.: Springer, Berlin, Heidelberg, 2011, vol. 6848, pp. 250-325, doi: 10.1007/978-3-642-23032-5_5 .

[48] S. Jupp, S. Bechhofer, and R. Stevens, "A Flexible API and Editor for SKOS," in *The Semantic Web: Research and Applications (Lecture Notes in Computer Science)*, L. Aroyo et al., Eds.: Springer, Berlin, Heidelberg, 2009, vol. 5554, pp. 506-520, doi: 10.1007/978-3-642-02121-3_38 .

[49] T. Tudorache, N. F. Noy, S. Tu, and M. A. Musen, "Supporting Collaborative Ontology Development in Protégé," in *The Semantic Web - ISWC 2008 (Lecture Notes in Computer Science)*, A. Sheth et al., Eds.: Springer, Berlin, Heidelberg, 2008, vol. 5318, pp. 17-32, doi: 10.1007/978-3-540-88564-1_2 .

[50] L. Halilaj, I. Grangel-González, G. Coskun, S. Lohmann, and S. Auer, "Git4Voc: Collaborative Vocabulary Development Based on Git," *International Journal of Semantic Computing*, vol. 10, no. 2, pp. 167-191, June 2016, doi: 10.1142/S1793351X16400067 .

[51] VoCol - An Integrated Environment for Colloborative Vocabulary Development. [Online]. Available: http://vocol.iais.fraunhofer.de/

[52] M. Salvadores, P. R. Alexander, M. A. Musen, and N. F. Noy, "BioPortal as a dataset of linked biomedical ontologies and terminologies in RDF," *Semantic Web*, vol. 4, no. 3, pp. 277-284, 2013, doi: 10.3233/SW-2012-0086 .

[53] G. Mochón, E. M. Méndez, and G. Bueno de la Fuente, "27 pawns ready for action: A multi-indicator methodology and evaluation of thesaurus management tools from a LOD perspective," *Library Hi Tech*, vol. 35, no. 1, pp. 99-119, 2017, doi: 10.1108/LHT-11-2016-0123 .

[54] A. Gonzales-Aguilar, M. Ramírez-Posada, and D. Ferreyra, "TemaTres: software para gestionar tesauros," *El profesional de la información*, vol. 21, no. 3, pp. 319-325, 2012, doi: 10.3145/epi.2012.may.14 .

[55] C. Mader, B. Haslhofer, and A. Isaac, "Finding Quality Issues in SKOS Vocabularies," in *Theory and Practice of Digital Libraries (Lecture Notes in Computer Science)*.: Springer, Berlin, Heidelberg, 2012, vol. 7489, pp. 222-233, doi: 10.1007/978-3-642-33290-6_25 .

[56] M. Fiorelli, M. T. Pazienza, S. Petruzza, A. Stellato, and A. Turbati, "Computer-aided Ontology Development: an integrated environment," in *New Challenges for NLP Frameworks 2010 (held jointly with LREC2010)*, La Valletta, Malta, 2010, 22 May, 2010. [Online]. Available: http://www.lrec-conf.org/proceedings/lrec2010/workshops/W10.pdf#page=33

[57] M. Fiorelli, R. Gambella, M. T. Pazienza, A. Stellato, and A. Turbati, "Semi-automatic Knowledge Acquisition through CODA," in *Modern Advances in Applied Intelligence - 27th International Conference on Industrial Engineering and Other Applications of Applied Intelligent System, IEA/AIE 2014*, vol. Part II, Kaohsiung, Taiwan, 2014, pp. 78--87, doi: 10.1007/978-3-319-07467-2_9 , June 3-6.

[58] World Wide Web Consortium (W3C). (2017, July) World Wide Web Consortium (W3C). [Online]. Available: https://www.w3.org/TR/shacl/

[59] J. McCrae, E. Montiel-Ponsoda, and P. Cimiano, "Collaborative semantic editing of linked data lexica," in *Proceedings of the Eight International Conference on Language Resources and Evaluation (LREC'12). Istanbul Lüfti Kirdar Convention & Exhibition Centre, Turkey, 21-27 May 2012*, 2012, pp. 2619-2625. [Online]. Available: http://www.lrec-conf.org/proceedings/lrec2012/pdf/544_Paper.pdf

[60] M. Rico and C. Unger, "Lemonade: A Web Assistant for Creating and Debugging Ontology Lexica," in *Natural Language Processing and Information Systems (Lecture Notes in Computer Science).*: Springer, Cham, 2015, vol. 9103, pp. 448-452, doi: 10.1007/978-3-319-19581-0_45 .

[61] A. Bellandi, E. Giovannetti, S. Piccini, and A. Weingart, "Developing LexO: A Collaborative Editor of Multilingual Lexica and Termino-ontological Resources in the Humanities," in *Proceedings of Language, Ontology, Terminology and Knowledge Structures Workshop (LOTKS 2017),co-located with the 12th International Conference on Computational Semantics (IWCS), 19 September 2017 Montpellier*, 2017. [Online]. Available: http://www.aclweb.org/anthology/W17-7010

[62] A. Bellandi, E. Giovannetti, and A. Weingart, "Multilingual and Multiword Phenomena in a lemon Old Occitan Medico-Botanical Lexicon," *Information*, vol. 9, no. 3, 2018, doi: 10.3390/info9030052 .

[63] F. Bond and K. Paik, "A survey of wordnets and their licenses," in *Proceedings of the 6th Global WordNet Conference (GWC 2012). Matsue, Japan, January, 9-13, 2012* , pp. 64-71.