BiOSS: A system for biomedical ontology selection

Marcos Martínez-Romero (corresponding author)
IMEDIR Center
Universidade da Coruña (UDC)
Campus de Elviña s/n, 15071, A Coruña, Spain
marcosmartinez@udc.es
Phone: +34-981167000-ext. 2667 / Fax: 981167172

José M. Vázquez-Naya
Departamento de Tecnologías de la Información y las Comunicaciones
Facultad de Informática (FIC)
Universidade da Coruña (UDC)
15071, A Coruña, Spain
jose@udc.es

Javier Pereira
IMEDIR Center
Universidade da Coruña (UDC)
Campus de Elviña s/n, 15071, A Coruña, Spain
javierp@udc.es

Alejandro Pazos
Departamento de Tecnologías de la Información y las Comunicaciones
Facultad de Informática (FIC)
Universidade da Coruña (UDC)
15071, A Coruña, Spain
apazos@udc.es

Abstract

In biomedical informatics, ontologies are considered a key technology for annotating, retrieving and sharing the huge volume of publicly available data. Due to the increasing amount, complexity and variety of existing biomedical ontologies, choosing the ones to be used in a semantic annotation problem or to design a specific application is a difficult task. As a consequence, the design of approaches and tools addressed to facilitate the selection of biomedical ontologies is becoming a priority. In this paper we present BiOSS, a novel system for the selection of biomedical ontologies. BiOSS evaluates the adequacy of an ontology to a given domain according to three different criteria: (1) The extent to which the ontology covers the domain; (2) The semantic richness of the ontology in the domain; (3) The popularity of the ontology in the biomedical community. BiOSS has been applied to 5 representative problems of ontology selection. It also has been compared to existing methods and tools. Results are promising and show the usefulness of BiOSS to solve real-world ontology selection problems. BiOSS is openly available both as a web tool and a web service.

1 INTRODUCTION

In recent years, ontologies have become an increasingly important topic in the biomedical field, where they are considered a critical component of multidisciplinary and translational research [1, 2]. The necessity of structuring, sharing and reusing the huge volume of data about genes, proteins, diseases, molecular functions, etc. that has been generated during the last years has led to the development of multiple biomedical ontologies, which are usually stored into large-scale repositories available for researchers. One of the best-known repositories of biomedical ontologies is NCBO's BioPortal [3], which is a Web-based, open repository containing more than 350 biomedical ontologies and terminologies, and this number is continuously growing.

Several authors have pointed out that reusing existing ontologies rather than creating new ones is a desirable practice [4-7]. Building an ontology from scratch is a very complex and time-consuming process, which requires specialized human resources. Reusing ontologies minimizes the time and cost of development and avoids the existence of multiple ontologies that represent the same meaning, which is necessary to ensure proper interoperability.

However, there are some barriers that hinder ontology reusing. As an example, suppose that a researcher needs to find an ontology containing the concept "Human Immunodeficiency Virus". The first step would be to search for the subset of current biomedical ontologies that contain such notion. He/she would find multiple sources, developed by different organizations and with different size and level of detail, or granularity (e.g. NCI Thesaurus, Medical Subject Headings, CRISP Thesaurus, etc.). Then, the second step would be to choose the most appropriate ontology for his/her specific problem. Achieving this process as a whole by hand is a difficult task that would require previous experience in ontology search and knowledge about current biomedical ontologies. As another example, suppose that the researcher needs to find an ontology containing not just a simple biomedical notion, but a set of 100 different concepts. Ideally, he/she would expect to find a single ontology containing all of them, but in practice this is very seldom the case. The concepts will usually be distributed over multiple ontologies. In general, the researcher would like to obtain the minimal set of ontologies that cover the maximal

set of concepts, and if two ontologies cover the same concepts, it will be necessary to select the best one. Due to the high number, size and complexity of current biomedical ontologies, carrying out this task manually is very complex and time-consuming.

This paper makes the following contributions: (i) We present BiOSS (Biomedical Ontology Selection System), which is an open access system that helps to choose the right biomedical ontology for a given domain; (ii) We compare BiOSS to the most relevant ontology search and selection approaches in biomedicine; (iii) We conduct an expert-based evaluation of BiOSS in several representative ontology selection scenarios. Our research is relevant both for researchers who need to select the best biomedical ontology for their domain and for Semantic Web developers who want to provide their systems with automatic biomedical ontology selection capabilities.

The rest of the paper is organized as follows. In section 2, related work is reviewed and BiOSS is compared with existing approaches. Section 3 presents the high-level system's workflow and describes the ontology selection process step by step. Section 4 provides implementation details and outlines the functionality of the system from a user perspective. Section 5 shows and discusses the results of evaluating the system in several knowledge reusing scenarios. Finally, section 6 concludes the paper and sketches some future lines of interest.

2   COMPARISON WITH RELATED WORK

Ontology selection is complementary to the process of ontology evaluation. Sabou et al. defined ontology selection as "the process that allows identifying one or more ontologies or ontology modules that satisfy certain criteria" and stated that this process is, "in essence, an ontology evaluation task" [8]. First contributions in the field of ontology evaluation date back to the early 1990s, starting with the work done by Gruber and Gómez-Pérez at the Knowledge System Laboratory (KSL) of Stanford University [9-11]. These works were motivated by the necessity of having evaluation strategies to guide and improve the ontology building process.

Some years later, with the birth of the Semantic Web [12], the need for evaluation methods addressed to choose the most appropriate ontologies for specific domains or applications emerged. This resulted in a considerable number and variety of methods and systems, aimed to retrieve all ontologies from a repository satisfying some basic requirements (e.g. find all ontologies that contain the keyword "cardiovascular disease"). In general, these approaches were restricted to one keyword search and were classified under the "ontology search" category.

Nevertheless, ontology selection goes a step beyond ontology search. Ontology selection implies not just to obtain a list of ontologies that contain class names matching a specific term, but to evaluate all candidate ontologies according to a set of criteria and provide a selection about the best ones for a specific domain or application [13, 14]. From this strict point of view, ontology selection is a complex process that comprises taking into account multiple input keywords, semantically expand them, evaluating the candidate ontologies according to a variety of criteria and providing a list of ontologies ranked according to their suitability to semantically represent the given keywords. For a review of the ontology evaluation and selection methods that have been proposed up to date, the reader might refer to [8, 15-18].

In spite of all the work that has been done during the last two decades, the application of ontology evaluation techniques to the biomedical field is still recent. This fact contrasts sharply with the large number of biomedical ontologies that exist today, and with the growing need for strategies and tools to select the best ontologies for describing biomedical resources.

Inspired by [13, 16], we have compared BiOSS with the 6 most relevant contributions for ontology search and selection in biomedicine: the ontology search mechanism provided by BioPortal[1]; the EBI's Ontology Lookup Service (OLS)[2]; the ontology selection approach conceived by Tan and Lambrix [19]; the ontology selection tool proposed by Maiga [20]; the

_____

[1] http://bioportal.bioontology.org/

biomedical ontology-search technique developed by Alani et al. [21]; and the NCBO's recommender service [16]. We have reviewed these approaches according to 18 criteria, grouped into 6 categories (see Table 1). These criteria are described below.

Table 1. Comparison of existing ontology selection approaches for the biomedical domain. A hyphen ('-') means undetermined or not applicable.

Feature / Method or tool		BioPortal Search	EBI OLS	Tan & Lambrix	Maiga	Alani et al.	NCBO Recommender	BiOSS
General	Type of method	search	search	selection	selection	selection	selection	selection
	Availability	tool, service	tool, service	no	no	no	tool, service	tool, service
	Automation	yes	yes	-	no	yes	yes	yes
	Approximate execution time (1/10/50 terms)	5s / - / -	5s / - / -	-	-	-	5s/10s/10s	7s/30s/4m
Candidate ontologies	Repository size	372	82	-	-	-	339	200
	Subset	yes	one ontology	-	no	-	yes	yes
	External ontologies	yes	no	-	no	-	yes	no
Input	Type of input	single keyword	single keyword	text	task and requirements	keywords	text, keywords	keywords
	Weighted keywords	-	-	-	-	no	no	no
Preprocessing	Spell checking	no	no	-	-	yes	no	yes
	Disambiguation	no	no	no	no	no	no	yes
	Query expansion	no	no	-	no	yes	no	yes
Selection criteria	Input coverage	yes	yes	yes	yes	yes	yes	yes
	Knowledge richness	no	no	-	yes	no	no	yes
	Knowledge formality	no	no	-	no	no	no	no
	Popularity	no	no	-	no	no	no	yes
Output	Single ontology	yes	yes	yes	yes	yes	yes	yes
	Combinations	-	-	no	no	no	no	yes

2.1 General criteria

- *Type of method.* This criterion refers to whether the method consists in identifying the list of biomedical ontologies that contain a given term (ontology search), or if it is a more complex process that takes into account several evaluation criteria and provides a selection of the best ontologies for a specific situation (ontology selection). BioPortal Search and EBI OLS are ────── ontology search approaches, while the others are considered ontology selection methods.

[2] http://www.ebi.ac.uk/ontology-lookup/

- *Availability*. This can be considered a critical issue when analyzing contributions addressed to solve real problems. Ideally, the approach should be available for users either as a software tool, so that it could be used manually, or as a web service, which would allow its integration into automatic knowledge reusing processes. As Table 1 shows, only 4 of the 7 approaches that have been reviewed are available for users. Moreover, only 2 of these approaches (The NCBO Recommender and BiOSS) can be considered pure ontology selection approaches.

- *Automation*. Automatic knowledge reuse on the Semantic Web requires fully automatic selection approaches that do not require user intervention during the selection process. The automation criterion indicates if the tool is able to achieve ontology selection in a fully automatic way or not.

- *Execution time*. As explained in [13], it is important to maintain a balance between the complexity of the evaluation methods and their performance, because ontology selection services should be fast enough to be dynamically invoked by client applications. Table 1 shows the reference execution times in seconds (s) or minutes (m) for each method given 1, 10 and 50 input terms. It can be seen that the NCBO Recommender is considerably faster than BiOSS when working with big sets of input terms. This is because BiOSS achieves a complex preprocessing step that even though is useful to improve the quality of the results obtained, it slows down the selection process. More details about the execution time of BiOSS will be provided in section 5.3.1.

2.2   Candidate ontologies

- *Repository size*. The number of candidate ontologies is another issue that requires attention. Ideally, a selection tool should take into account all, or at least most part of existing biomedical ontologies. The current version of BiOSS includes 200 of the most relevant and widely accepted

biomedical ontologies from several well-known resources such as the Unified Medical Language System (UMLS) [22], BioPortal[3] and the Open Biomedical Ontologies repository[4].

- *Subset.* When a user needs to find an ontology for a specific domain (e.g. mouse anatomy, infectious diseases, etc.), he/she may want to limit the selection process to those ontologies most relevant to such domain, instead of using all candidate ontologies provided by the system. The "subset" criterion refers to the possibility of restricting the ontology selection process to a specific subset of the candidate ontologies.

- *External ontologies.* This criterion refers to whether the system is restricted to a particular list of candidate ontologies or it allows that ontologies uploaded by the user can participate in the selection. The current version of BiOSS does not permit users to upload external ontologies. The selection process is restricted to a fixed ontology repository that is periodically updated by the system's administrators.

2.3 Input

- *Type of input.* Ontology selection requires as input some kind of representation of the specific domain in order to be able to look for the most adequate ontologies for it. Most basic approaches restrict the query to a single keyword (e.g. "anatomy"). This type of input allows to decrease the number of candidate ontologies quickly and easily and is appropriate for basic searches such as ontology names (e.g. NCI Thesaurus) or well-known fields (e.g. cancer, infectious diseases, mental health, etc.). Nevertheless, a single keyword is not enough when it is necessary to face more specific problems such as finding the best ontology for annotating a biomedical corpus or to represent a specific set of terms. As a consequence, more advanced

_____

[3] http://bioportal.bioontology.org/
[4] http://www.obofoundry.org/

selection methods rely on a set of keywords or on a body of text from which a set of keywords is extracted.

- *Weighted keywords*. When providing a text or a set of terms as input, it is possible that not all keywords are equal to the user. Some terms could be more important to the searcher than others. One keyword may be more central to the domain or application and another may be linked to an elusive feature. This criterion refers to whether the tool allows to assign different importance (or weight) to the input keywords. The use of queries based on weighted keywords is a technique commonly used by search engines but none of the ontology selection methods developed so far has this feature.

2.4 Preprocessing

- *Spell checking*. The input text or keywords could contain mistakes (e.g. spelling errors, different case letters, punctuation symbols, double spaces, etc.) that must be corrected to guarantee a proper ontology selection. A selection tool aimed to solve real problems should be able to ensure that the selection process starts from error-free keywords. The spell checking functionality provided by BiOSS is based on Google's spell checking service.

- *Disambiguation*. Some input terms provided to the selection system may have more than one meaning. For example, the term "cold" may refer to a low temperature or to a common infection. This criterion refers to whether the ontology selection process takes into account any kind of mechanism to identify the intended meaning of the input terms that are ambiguous. As stated in [13], current ontology selection techniques treat the comparison between query terms and ontology concepts rather superficially, relying only on lexical comparisons. The matching methods do not take into consideration the semantics of the ontologies for concept recognition, which is considered as a major drawback of ontology selection approaches [16]. Up to our knowledge, BiOSS is the first biomedical ontology selection system that checks the semantic relations between the query terms and the ontology concepts.

9

- *Query expansion*. This is the process of supplementing the original query with semantically related terms (e.g. synonyms, hyperonyms, hyponyms, etc.) in order to achieve a better match between the system's input and the candidate ontologies. BiOSS uses UMLS to supplement the query terms with synonyms, in order to broaden the search space for ontologies when no ontologies are found for a term.

2.5   Selection criteria

- *Input coverage*. The goal of ontology selection is retrieving the ontology, or ontologies, that best represent the user's domain. As a consequence, measuring how well each candidate ontology covers the system's input is an obvious but fundamental step that has been taken into account by all the approaches that have been reviewed.

- *Knowledge richness*. Ontologies with a higher degree of detail in the representation of their concepts can be considered potentially more useful to describe a given domain that more simple ones. A good representation of a specific concept may include different types of relations to other concepts (e.g. is_a, part_of, etc.), an explanation about the concept, several concept attributes, etc. [23, 24]. The knowledge richness criterion refers to whether the richness of the conceptualization offered by the ontology is taken into account when achieving the evaluation.

- *Knowledge formality*. Existing biomedical knowledge sources present different levels of formality, from simple vocabularies and thesaurus to formal ontologies. This criterion reflects if the ontology selection method takes into account the formality of the ontology during the ontology selection process.

- *Popularity*. It refers to whether the selection method takes into account in some manner how well known and well established, or "popular", the ontology is in the biomedical community. Several approaches have already been proposed for the problem of measuring ontology popularity [24-26], which rely on the assumption that the relevance of an ontology is proportional to the number of ontologies that reference it, using the same principle as most search engines. Nevertheless, these methods are insufficient because current ontologies are

poorly connected and many of them are not referred to by any other ontologies at all [23]. Among the approaches that have been compared, BiOSS is the only one that takes into account ontology popularity.

2.6 Output

- Single ontology. This is the most common type of output and is present in all the approaches that have been studied. The system provides a list of ontologies, ranked according to the result of the evaluation process. Each position in the ranking corresponds to one and only one ontology.

- Combinations. This criterion refers to the capability of the approach to return possible combinations of ontologies that jointly provide a good representation of the user's domain. In this manner, each position in the ranking could be constituted by one or several ontologies. This criterion has been identified as one of the requirements of ontology selection by Sabou et al. [8] Among the approaches that have been compared, BiOSS is the only one capable of providing ontology combinations as output.

3 DESCRIPTION OF THE ONTOLOGY SELECTION APPROACH

BiOSS is based on a multi-criteria selection approach aimed to evaluate, in the shortest time possible, several key aspects of each biomedical ontology and its suitability for describing the input keywords. This section provides an overall description of the algorithm. For a comprehensive explanation and details about the methods used, the reader is referred to our previous work [27, 28].
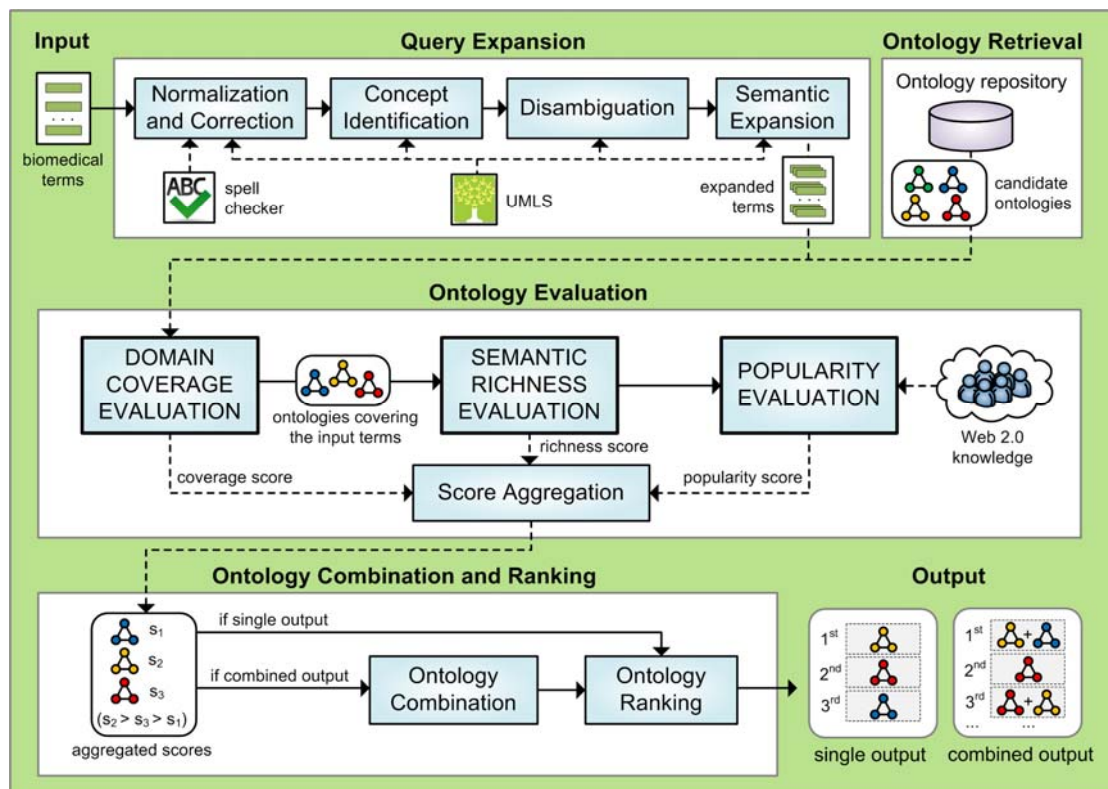
Fig. 1. Ontology selection workflow.

The ontology selection workflow is presented in Fig. 1. It shows the main components of the system, as well as the steps involved in the ontology selection process. The execution starts from a set of biomedical input terms that may be provided by a user, a Semantic Web agent, or automatically extracted from an information resource (e.g. text, image, webpage, etc.). Then, ontology selection is accomplished according to the following stages.

3.1 Query expansion

The purpose of this stage is to expand each initial term with other terms with the same meaning, or synonyms. These synonyms will be useful later, during the ontology evaluation process, to assess to what extent the initial terms are covered by each candidate ontology. The semantic expansion stage is supported by a widely known repository of biomedical terminology: the

Unified Medical Language System (UMLS) [22], developed by the US National Library of Medicine (NLM). Semantic expansion is achieved as follows:

Normalization and correction. Firstly, the initial terms are normalized to a common format by means of the "norm" tool, included in the UMLS Lexical Tools[5]. Normalization involves stripping possessives, replacing punctuation with spaces, removing stop words, lower-casing each word, breaking a string into its constituent words and sorting the words in alphabetic order. Then, a spell checker (Google's spell checker service, invoked through the Java API for Google spell checking[6]) is used to correct potential typographical mistakes.

Concept identification. Each input term may have one or several meanings. At this point, UMLS is queried to extract all the possible senses for each input term. Table 2 shows the possible meanings in UMLS for the terms "aorta" and "diaphragm". The term "diaphragm" has more than one possible sense, so the selection approach will have to choose one of them. This process is known as disambiguation, and it is explained below.

––––––

[5] http://www.ncbi.nlm.nih.gov/books/NBK9680/
[6] http://code.google.com/p/google-api-spelling-java/

Table 2. Example of input terms, their possible meanings in UMLS and their semantic types. The code that is displayed before the definition of each term (e.g. C0003483) corresponds to the UMLS concept unique identifier (CUI).

Input term	UMLS concept	UMLS semantic type
aorta	C0003483: The major arterial trunk that carries oxygenated blood from the left ventricle into the ascending aorta behind the heart, the aortic arch, through the thorax as the descending aorta and through the abdomen as the abdominal aorta (...)	Body Part, Organ, or Organ Component
diaphragm	C0011980: The musculofibrous partition that separates the thoracic cavity from the abdominal cavity. Contraction of the diaphragm increases the volume of the thoracic cavity aiding inhalation (...)	Body Part, Organ, or Organ Component
	C0042241: A medical contraceptive device of soft flexible material, usually of thin rubber, that is designed to cover the cervix uteri prior to sexual intercourse to prevent the entry of spermatozoa.	Medical Device

Disambiguation. In order to achieve a correct semantic expansion, each input term must be associated with a single meaning. In the case of ambiguous terms (terms with multiple senses), it is necessary to determine the most appropriate meaning according to the particular ontology selection scenario. Our word sense disambiguation (WSD) mechanism is based on the idea proposed in [29] of using the UMLS Semantic Network as a reference resource to disambiguate. The Semantic Network provides a categorization of medical and biomedical concepts into semantic types. In addition, it includes relations among semantic types. Our method consists in determining the semantic type that best represents the set of input terms (we call it "preferred semantic type"). After that, the disambiguation for each ambiguous term is achieved by selecting the meaning that is closer to the preferred semantic type in the Semantic Network, considering it as a graph. As a basic example, the UMLS semantic type that best represents the terms in Table 2 would be "Body Part, Organ, or Organ Component" because both terms belong to it. Then, the approach would consider that the most appropriate meaning for the term "diaphragm" is C0011980 because its semantic type is closer to the preferred semantic type than the other meaning (in fact, it is the same). Although it is not appropriate to use semantic types for disambiguation purposes in situations where the concepts linked to the ambiguous word are

assigned the same semantic type, a recent study has been demonstrated that this strategy achieves better results compared to other existing methods [30]. For a more comprehensive description of the disambiguation method the reader is referred to [31].

Semantic expansion. After disambiguation, each input term is associated to a single meaning. Then, UMLS is used to expand each term with its synonyms. These synonyms will be useful later, during the ontology evaluation stage, to increase the chances of finding each term in the candidate ontologies. As an example, the result of the semantic expansion for the term "leucocyte" would be the set of terms {"leukocyte", "leucocyte", "white blood cell", "white cell", "white blood corpuscle", "white corpuscle", "wbc"}.

As a result of the query expansion stage, several sets of synonyms are obtained (one for each input term), which can be considered a rich and error-free representation of the particular domain for which it is necessary to find ontologies. On the basis of this information, the candidate ontologies are evaluated.

3.2   Ontology retrieval

BiOSS has a local repository that stores all the ontologies that are susceptible to participate in the selection process. This repository is maintained and periodically updated by the system's administrators. The ontology retrieval stage involves accessing the repository and obtaining the set of candidate ontologies.

At the moment, great efforts are being made to develop large scale, openly accessible repositories of biomedical ontologies (e.g. the OBO Foundry initiative [32]). The current version of the system is based on a local ontology repository to avoid potential availability and performance issues caused by using a remote repository. However, the system's architecture does not impose any restriction with respect to the public or private character of the repository. BiOSS can be easily adapted to retrieve ontologies from a remote repository of biomedical ontologies that provides some way of accessing its content, like the NCBO's BioPortal [3].

3.3 Ontology evaluation

This stage constitutes the core of the ontology selection process. At this point, each candidate ontology is evaluated according to three different criteria: (1) **domain coverage**, or how well each candidate ontology covers the input terms; (2) **semantic richness**, which represents the degree of detail in the representation of the knowledge provided by each ontology; and (3) **popularity**, or relevance of the ontology in the biomedical community. In the following a general description of these criteria is provided. For a more detailed explanation, formal definitions and examples of how these criteria are calculated the reader is referred to our previous work [28].

Domain coverage. As previously explained, BiOSS is aimed to select the best ontology (or ontologies) to describe a set of input terms. The aim of the domain coverage evaluation is to assess to what extent the input terms are described by each candidate ontology or, in other words, to find the semantic mappings between the input terms and the ontology concepts. It consists in calculating the percentage of input terms that are contained in each ontology, considering that a term is contained in the ontology if there is an ontology class whose name matches the term (after normalization and correction) or one of its synonyms. We only have considered exact matches because it has been found that taking into account partial matches may reduce the matching quality [13, 33]. While this mechanism may return some useful hits (e.g. "cardiologist" when searching for "cardiology"), it may also be a source of significant mistakes (e.g. "chancellor" or "excellent" when searching for "cell"). The quality of the mappings identified has a high impact on the overall performance of the system because an ontology may contain thousands of concepts, be semantically rich and very popular in the research community, but if it does not provide a good coverage of input terms, it will not be useful to describe the domain. Due to this, ontologies that do not provide a minimum coverage are discarded at this point. The current version of BiOSS discards the candidate ontologies that provide a coverage lower than 10%.

Semantic richness. Ontologies that offer more detail for a given domain can be considered potentially more useful than less detailed ontologies. When accessing an ontological concept (e.g. "endocardium"), it is valuable to encounter some additional information like a definition (e.g. "The layer of endothelial cells and connective tissue lining the chambers of the heart"), its direct relatives (e.g. "is_a" "heart part"), relations to other concepts (e.g. "part_of", "located_in", "contained_in", etc.). Semantic richness evaluation is aimed to measure the level of detail in the representation of the knowledge provided by the ontology for the input keywords.

The method followed by BiOSS to evaluate the semantic richness is independent of the language in which the ontology is expressed (e.g. OWL) and it only evaluates the richness of the concepts that are involved in the domain coverage, instead of taking into account the semantic richness of all the concepts in the ontology. It is based on evaluating the following factors:

- Concept relatives. The amount of relatives for each concept is one of the aspects that may be considered important when measuring the semantic richness of an ontology. BiOSS takes into account the amount of direct parents, siblings and children for each concept.

- Additional information. Apart from its relatives, an ontology typically contains other information for each concept. This step is addressed to measure this kind of information. BiOSS takes into account all characteristics of each concept that provide information about such concept, with the exception of: (i) its relatives, which were already assessed; (ii) the concept name (or label), because we assume that every concept will have a name, and we will not consider it a distinguishing feature; and (iii) the concept instances, because we believe that they do not necessarily reflect the richness of the conceptual structure itself. Having this into account, the additional information for a concept would consist on the number of relations with other concepts, definitions, synonyms, restrictions over values or datatypes, etc.

- Similar knowledge. In fields such as biomedicine, it is common to find ontologies that, besides containing a particular concept (e.g. "heart"), also provide many other concepts with similar semantics ("heart disease", "cardiac insufficiency", "heart process"). The evaluation of similar knowledge is addressed to measure how much knowledge contains an ontology related to a specific concept. It is based on counting the number of concepts in the ontology whose name contains the name of the specific concept or any of its UMLS synonyms (e.g. "heart process" contains "heart" so BiOSS considers that they are concepts with similar semantics).

Popularity. Apart from assessing how well each ontology covers a specific domain and the richness of its structure with respect to it, there is another aspect that requires special attention: what if the knowledge in the ontology is wrong? According to the definition of ontology provided by Studer [34], "an ontology captures consensual knowledge, that is, it is not private to some individual, but accepted by a group". With this in mind, any method addressed to evaluate an ontology should also take into account the level of collective acceptability or popularity of the ontology.

BiOSS uses an innovative method to quantify the popularity of an ontology. It is based on the collective knowledge stored in widely-known Web 2.0 resources, whose value is created by the aggregation of many individual user contributions. It consists of counting the references to the ontology from each Web 2.0 resource. Depending on the kind of resource, the way of measuring the references will be different.

The current version of BiOSS uses the following resources to measure ontology popularity:

- PubMed. Ontologies referenced in scientific papers are generally well known, mature and trustworthy ontologies. For each candidate ontology, BiOSS counts the number of papers in PubMed in which the name of the ontology appears at least once.

- BioPortal[7]. It is the reference resource for publishing biomedical ontologies. For each ontology, the system checks if it is indexed in BioPortal or not.

- Wikipedia. In a similar way than for PubMed, BiOSS counts the number of articles in Wikipedia that contain at least once the name of the ontology.

- Twitter. In this resource, the number of references to the ontology are calculated as the number of tweets that contain at least once the name of the ontology.

Each kind of evaluation (i.e. domain coverage, semantic richness and popularity) provides a numerical score as a result. The three resulting scores are aggregated into a single value that indicates the overall goodness of the ontology to represent the given input terms. This aggregation required to determine the set of weights that allowed to give more or less importance to each parameter. This adjustment was based on the opinion of 5 experts in biomedical ontologies, which have expressed their view about the values to be used by filling a short questionnaire (available at http://tinyurl.com/oe57k8a). According to the results obtained, the BiOSS web interface has been configured with the weights shown in Table 3. In addition, the BiOSS web service interface allows the user to invoke the selection process with the weights he/she considers most appropriate. For a formal definition and a comprehensive explanation of how all these scores are calculated the reader is referred to [28].

---

[7] http://bioportal.bioontology.org/

Table 3. Weights used to configure the Web user interface.

Criterion/Parameter		Weight
	Domain coverage	0.55
Final score	Semantic richness	0.26
	Popularity	0.19
	Relatives	0.33
Semantic richness	Similar knowledge	0.34
	Additional information	0.33
	Wikipedia	0.26
	Twitter	0.16
Popularity	BioPortal	0.28
	PubMed	0.30

3.4 Combination and ranking

BiOSS provides two different kind of output: "single" or "combined". The single output consists in ranking the candidate ontologies according to the aggregated score obtained from the evaluation process. The higher the ontology is located in the ranking, the better it will be to describe the initial terms. In the combined output, each position in the ranking is not just a single ontology. It may be a set of ontologies, that can be used together to provide a better coverage of the domain. In this case, the aggregated score of each combination of ontologies is calculated from the aggregated scores obtained individually for each ontology. For each combination of ontologies, the following parameters are considered:

- Combined domain coverage. It is calculated taking into account the terms covered by the combination of ontologies as a group. If there are some terms covered by more than one ontology, the ontology with the highest semantic richness value is chosen.

- Combined semantic richness and popularity. They are calculated on the basis of the semantic richness and popularity for each individual ontology, proportionally to the number of terms covered by each ontology.

The final score for each combination of ontologies is calculated in the same way than for the single output but using these combined scores and finally these scores are ranked.

4   IMPLEMENTATION DETAILS

As explained by Cantador and colleagues [14], there are two well-known scenarios for ontology reuse. The first one refers to the situation where a user needs to find the best ontology to describe a specific domain. The second one is related to the case in which Semantic Web applications need to automatically and dynamically select ontologies.

BiOSS has been designed with both scenarios in mind. It has a web service interface, which makes it possible to invoke it in situations where automatic ontology reusing is required (e.g. automatic annotation of biomedical resources [35]). In addition, a web user interface has been developed on the basis of the web service, to make the system's functionality available to common users (see Fig. 2).
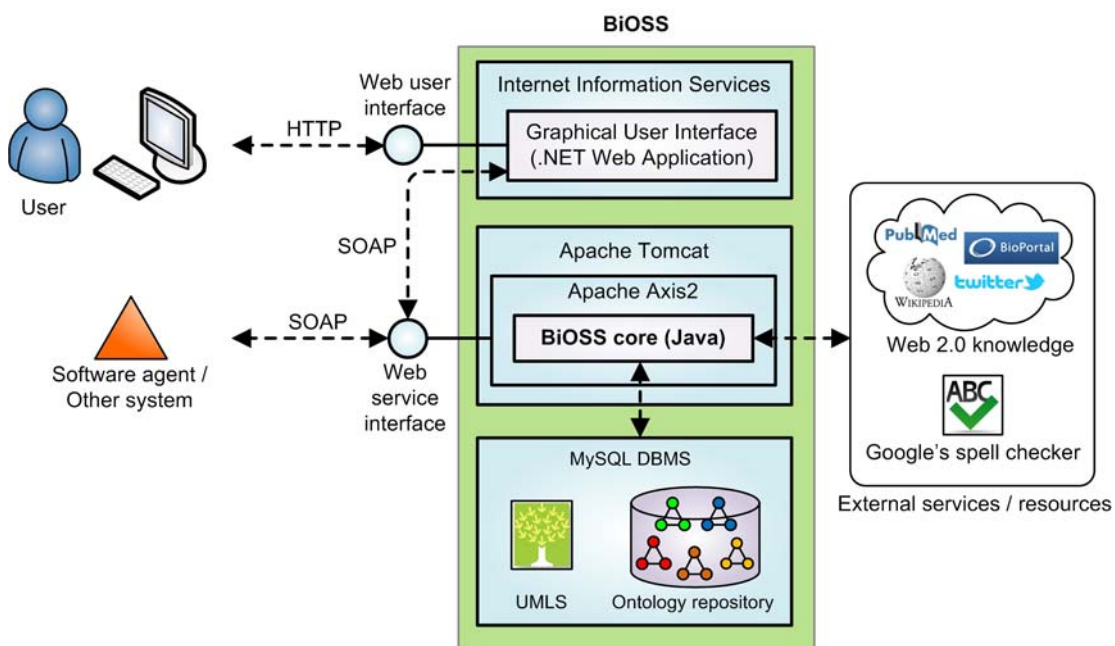

Fig. 2. System's overall architecture.

The web service interface[8] has been developed using Java technology by means of the Apache Axis 2 framework. It follows the SOAP protocol and has been deployed on the Apache Tomcat

_____

[8] http://193.147.41.219:8080/axis2/services/BiOSSWService?wsdl

container. The web user interface[9] has been developed using Microsoft .NET technology and interacts with the SOAP web service to provide an intuitive ontology selection functionality.

The web user interface is shown at Fig. 3. It has an input text field to manually enter the keywords representing the user's domain. Several predefined sets of input terms for specific fields (e.g. "Breast cancer", "Critical care", etc.) that were used to evaluate the system are provided as an example on the Settings-Input panel. Details about the necessity of selecting ontologies for each test case are shown by clicking the link "Description of test cases". By clicking the "Go!" button, the ontology selection process starts.
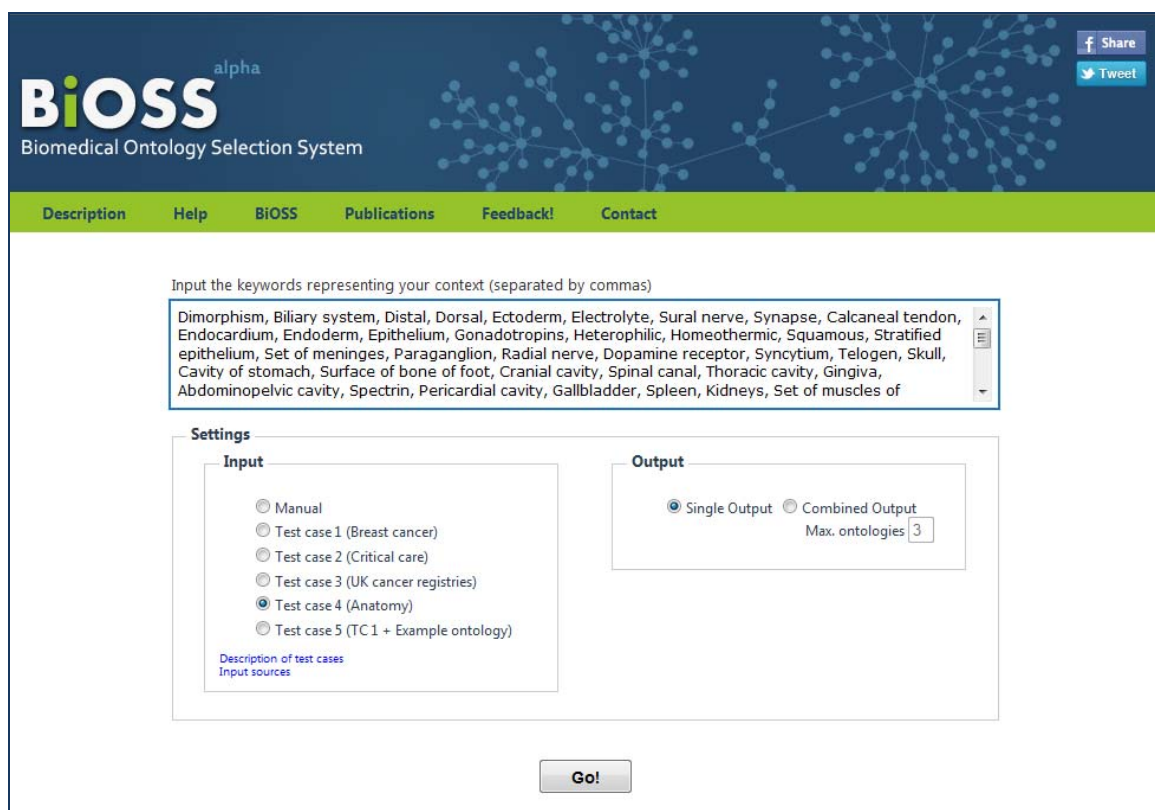

Fig. 3. Web user interface of BiOSS.

___

[9] http://bioss.ontologyselection.com

22

Currently, the system works on top of a repository of 200 biomedical ontologies. By clicking the "Input sources" link, the user can see the list of candidate ontologies and uncheck the ones that he/she prefers to exclude from the selection process.

The repository has been implemented as a MySQL database following a relational model that has been specifically designed to meet the requirements of our ontology selection approach (see Fig. 4). For each ontology, the repository stores all its concepts and several precomputed values that are needed to achieve the evaluation according to our approach (e.g. references from Web 2.0 resources such as Wikipedia and Twitter, number of relatives for each concept, etc.)
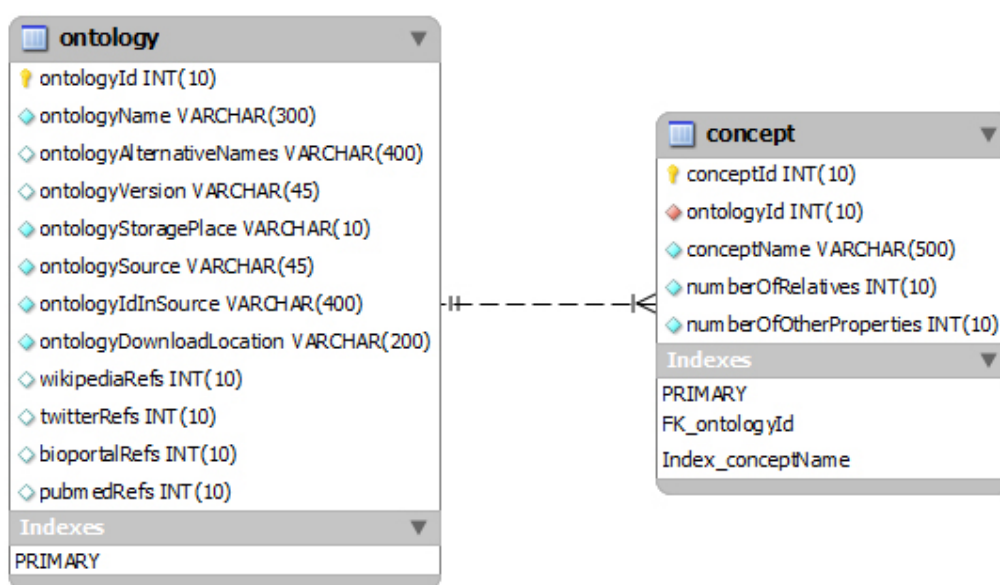

Fig. 4. Database model of the ontology repository.

With respect to the output, the system provides two alternatives: single or combined. When choosing the single output, the system will provide a ranked list of ontologies for the input terms. It is the preferred output in most cases. However, when it is necessary to find ontologies to describe the largest possible number of multiple input terms, it is not usual to find a single ontology containing all of them. In such situations, it would be useful to use two, three or even more ontologies to cover the largest possible number of input terms [13]. When selecting the combined output, the system evaluates how well different sets of ontologies cover the initial

terms in order to get the best results. Doing this task by hand would consume several hours or even days, but the BiOSS system is capable of providing an output in some seconds.

Fig. 5 shows a fragment of single output for a set of five input keywords: "white cell", "chemotherapy", "DNA", "apoptosis" and "cavity of stomach". The following information is shown:

- The **position** of the ontology in the ranking. Ontologies that appear higher in the ranking are supposed to be better to describe the initial terms than ontologies that appear lower. In this case, BiOSS recommends the user to select Medical Subject Headings (MeSH).

- The name of the **source** (ontology). Some examples are Medical Subject Headings, NCI Thesaurus, etc.

Position	Sources	Final score	Covered	Not covered	Coverage	Richness	Popularity
1	Medical Subject Headings	0.759	chemotherapy(pharmacother) | white cell(white blood corpuscles) | apoptosis(programmed cell death, type i)	cavity of stomach | DNA	60.0%	0.918	1.000
2	NCI Thesaurus	0.737	chemotherapy | white cell(wbc) | apoptosis(pcd)	cavity of stomach | DNA	60.0%	0.934	0.863
3	Logical Observation Identifier Names and Codes	0.709	chemotherapy | white cell(leukocytes) | DNA	cavity of stomach | apoptosis	60.0%	0.951	0.697

Fig. 5. Example of single output. Only the top three positions in the ranking are shown.

Position	Sources	Final score	Covered	Not covered	Coverage	Richness	Popularity
1	Medical Subject Headings Gene Ontology Foundational Model of Anatomy	0.951	chemotherapy(pharmacother) | white cell(white blood corpuscles) | cavity of stomach | apoptosis(programmed cell death, type i) | DNA(dna location)		100.0%	0.891	0.891
2	NCI Thesaurus Gene Ontology Foundational Model of Anatomy	0.938	chemotherapy | white cell(wbc) | cavity of stomach | apoptosis(pcd) | DNA(dna location)		100.0%	0.901	0.809
3	Medical Subject Headings Gene Ontology Cardiac Electrophysiology Ontology	0.935	chemotherapy(pharmacother) | white cell(white blood corpuscles) | cavity of stomach | apoptosis(programmed cell death, type i) | DNA(dna location)		100.0%	0.857	0.856

Fig. 6. Example of combined output. Only the top three positions in the ranking are shown.

- The **final score** obtained by the system for each ontology (range [0, 1]). The list of ontologies is ranked according to this score. This aggregated measure is calculated by combining the values of ontology coverage, ontology richness and ontology popularity. These values are shown in the three columns to the right of the ranking results, and will be discussed later.

- The initial terms that are **covered** and **not covered** by the ontology. An initial term is covered by the ontology if its meaning is contained in the ontology. If the ontology contains the meaning

24

of the term but using a synonym, then the concept name is shown between brackets. As an example, the text "white cell (white blood corpuscles)" at the first position in the ranking of Fig. 5, means that the input term "white cell" is covered by Medical Subject Headings through the concept named "white blood corpuscles".

- Coverage provided by the ontology for the input terms, that is, the amount of input terms contained in the ontology with respect to the total number of input terms. It is expressed as a percentage. A coverage of 100% means that all input terms are covered by the ontology.

- Ontology **richness**. It represents the level of detail provided by the ontology for each input term. It ranges from 0 to 1.

- Ontology **popularity**. Level of collective acceptability or popularity of the ontology according to its relevance in widely known Web 2.0 resources (e.g. Wikipedia, BioPortal, etc.). It ranges from 0 to 1.

Fig. 6 shows the combined output for the input terms of the previous example. When choosing this kind of output, the system is allowed to select more than one ontology to provide the maximum domain coverage, so each place in the ranking will be usually constituted by several ontologies instead of a single one.

When executing the single output (Fig. 5), the first ontology in the ranking (i.e. MeSH) provides a domain coverage of 60%. However, the combined output (Fig. 6) shows that it is possible to reach a coverage of 100% by means of using several ontologies together (i.e. MeSH, GO and FMA). In the combined output, each ontology covers some specific terms to reach the maximum coverage as a whole. Different colors are used to show the terms covered by each ontology. If we look at the first position in the ranking of Fig. 6, it is possible to interpret that GO covers the term "DNA", FMA covers "cavity of stomach" and MeSH covers the rest of the terms. The name of each ontology links to a new page that contains information about the ontology and a link to access or download it.

5 EVALUATION

The system has been applied to 5 scenarios (test cases) representative of the biomedical ontology selection field. We have contacted by e-mail several experts in biomedical ontologies and asked them to execute BiOSS in at least one of the mentioned scenarios, to analyze the results obtained and to fill out a questionnaire available online.

In all test cases, the expert evaluation has been limited to the first 10 results provided by BiOSS. Both the single and the combined output have been taken into account. For the combined output, the maximum number of ontologies to combine has been set to 5. With respect to the domain coverage evaluation, a threshold of 0.1 has been set. This means that ontologies that covered less than a 10% of the input keywords were directly discarded.

In the following, the test cases and the questionnaire used for the evaluation are described. After that, the evaluation results are presented.

5.1 Test cases

The system has been applied to the following ontology selection scenarios:

- **Test case 1: Breast Cancer (30 terms)**. Researchers at Stanford University are building a system that abstracts clinical information from two electronic medical record databases related to the care and management of breast cancer. To build this application, they need to reuse ontologies already developed by other organizations. This scenario has been previously used by Jonquet et al. to evaluate the NCBO's Ontology Recommender Web Service [16].

- **Test case 2: Critical Care (32 terms)**. Researchers at the University of A Coruña (Spain) are building an ontology related to the field of Cardiac Intensive Care Units (CICUs). This ontology will be used as the basis of an expert system for remote monitoring and control of patients who underwent heart surgery. In order to build such ontology, these researchers want to reuse knowledge from existing ontologies. A deeper explanation about this work can be found at [36].

- **Test case 3: UK Cancer Registries (24 terms)**. The UK is covered by 11 cancer registries, coordinated by the United Kingdom Association of Cancer Registries[10] (UKACR). These registries, which have stored population-based information on cancer from more than 40 years ago, contain a common set of 24 epidemiological variables. In order to facilitate the future integration of these data with data from other registries, in this test case we suppose that it is necessary to find the best ontologies to semantically annotate the mentioned variables.

- **Test case 4: Anatomy (80 terms)**. In this test case we assume that it is necessary to find the best ontology or ontologies to annotate a set of terms randomly chosen from the field of anatomy.

- **Test case 5: Synthetic Ontology (30 terms)**. The purpose of this test case is testing the capability of the system to discard ontologies that are not relevant for the biomedical community. We suppose the situation described in the test case 1 (Breast Cancer), but adding a synthetic ontology as a candidate, which we created specifically for this test case using the Protégé Ontology Editor[11]. This ontology contains 30 concepts that match the 30 initial terms in the test case 1 (100% coverage), but it is not a shared ontology and it is not known by the biomedical community.

The input terms used in each scenario can be queried at the BiOSS website[12].

5.2 Evaluation questionnaire

The questionnaire used for the evaluation is available online[13] and contains the following sections:

---

[10] http://www.ukacr.org/
[11] http://protege.stanford.edu/
[12] http://bioss.ontologyselection.com/TestCases.aspx
[13] http://tinyurl.com/5uobbnw

- **Personal information (PI).** Contact information about the evaluator (name, surname, institution, e-mail and country). This section is compulsory.

- **Test cases (5 sections: TC1, TC2, TC3, TC4, TC5).** These sections are aimed to evaluate the results provided by BiOSS in each test case. We know that filling out the questionnaire for the 5 test cases is a hard task that may take a considerable amount of time. Due to this, the evaluators were not asked to fill out all the test cases, but at least one of them. Each evaluator selected the best ontologies for a test case. He/she was allowed to use any tool or website that he/she wanted. After that, he/she was asked to contrast his/her selection with the results provided by the system. Both single and combined outputs were considered. Then, the evaluator was asked to assess the helpfulness of the system in each test case and give it a mark, according to the following options (test cases 1-4): Not helpful (0); Few helpful (1); Helpful (2); Very helpful (3). For the test case 5, the evaluators were asked to provide their opinion about the position of the "Synthetic Ontology" in the ranking, according to the following options: (1) The position is adequate; (2) The position is not adequate, it should be ranked higher; (3) The position is not adequate, it should be ranked lower; (4) I do not know if the position is adequate or not.

- **General questions (GQ).** This section contains several questions aimed to obtain the expert's opinion about the system, as well as about its main advantages and disadvantages. The evaluator was asked to provide a mark about the overall usefulness of the system, according to the following options: Not helpful (0); Few helpful (1); Helpful (2); Very helpful (3).

5.3 Evaluation results

In this section, we describe the results obtained after evaluating the system. Firstly, the execution times are presented. Finally, we summarize the feedback provided by the evaluators.

5.3.1   Execution times

Table 4 shows the execution time of BiOSS for each test case. These values were obtained with a PC computer (Intel Core 2 Quad 2.66GHz with 4GB of RAM) acting as a Web server. The time for each test case was calculated as the mean of five executions.

It is important to take into account that the execution time is strongly dependent of the amount of information about previous queries that is loaded into the server's memory (database cache). Table 4 differentiates between the time obtained in the worst situation (slowest execution) and the best situation (fastest execution). In the worst situation (e.g. after rebooting the server, or when none of the query terms have been searched since the last reboot), no previous queries are stored into memory, so BiOSS has to access the disk multiple times to obtain all the data necessary for ontology evaluation from the ontology database. In the best situation (e.g. the selection process is executed for the second time with the same input), all the data required by BiOSS are retrieved from the database cache and the access is much faster.

In order to measure the performance of BiOSS for an average scenario, we have analyzed the behavior of the database cache during the period July 1, 2013 – September 30, 2013 and have observed that for a total number of 468 executions made by 47 different users, the average of database cache hits has been 37.6%. Having this into account, we have configured the system to perform a 37.6% of accesses to the database cache and we have obtained the execution time for each test case (see Table 4).

Table 4. Average execution time for each test case (in seconds).

Test case	Terms	Candidate ontologies	Worst situation (100% accesses to DB cache)		Best situation (100% accesses to DB cache)		Average situation (37.6% accesses to DB cache)	
			Single output	Comb. output	Single output	Comb. output	Single output	Comb. output
TC1	30	200	353.27	354.15	3.87	3.92	207.21	208.70
TC2	32	200	181.31	182.18	7.03	8.08	109.26	109.72
TC3	24	200	121.06	121.75	4.15	6.95	75.17	75.89
TC4	80	200	621.92	631.15	7.28	22.34	398.37	409.12
TC5	30	201	364.65	364.87	3.90	3.96	221.07	221.82

In our opinion the execution times are acceptable for a general use of BiOSS, but these results could be improved to deal with problems that require to process large sets of terms in real time. These execution times could be enhanced in different ways, such as using a different storage approach, using a faster machine and optimizing the database cache.

5.3.2   Expert evaluation

As explained before, in order to evaluate the system we asked some experts to use it in at least one of the proposed scenarios and to fill out an online questionnaire. We have received feedback from 16 experts (see Table 5). Twelve of them (E1 to E12) have filled out at least one test case or the general questions section. The 4 remaining evaluators (E13 to E16) have not filled out the questionnaire. Instead, they have provided their opinion about the system and some suggestions to improve it by e-mail.

Table 5. Information provided by each evaluator.

Evaluator	Questionnaire sections completed	Additional comments
E1	PI, TC1, TC2, TC3, TC4, TC5, GQ	Yes
E2	PI, TC1, TC2, TC3, TC4, TC5, GQ	Yes
E3	PI, TC1, TC2, TC3, TC4, TC5, GQ	Yes
E4	PI, TC1, TC2, TC3, TC4, TC5, GQ	Yes
E5	PI, TC1, TC2, TC3, TC5, GQ	Yes
E6	PI, TC1, TC2, GQ	Yes
E7	PI, TC1, TC5, GQ	Yes
E8	PI, TC1, TC5, GQ	Yes
E9	PI, TC1, GQ	Yes
E10	PI, TC1, GQ	No
E11	PI, GQ	Yes
E12	PI, GQ	No
E13	None	Yes
E14	None	Yes
E15	None	Yes
E16	None	Yes

Table 6. Average of scores given by the evaluators for each test case (except for the test case 5, whose evaluation has been achieved differently) and regarding the overall usefulness of the system (GQ). The column "Evaluations received" refers to the number of evaluators that have filled out the corresponding section of the questionnaire. 0: Not helpful; 1: Few helpful: 2: Helpful; 3: Very helpful.

Questionnaire section	Evaluations received	Average score	
		Single output	Comb. output
TC1	10	2.30	2.00
TC2	6	2.33	2.17
TC3	5	2.40	2.20
TC4	4	2.75	2.00
GQ	12	2.50	2.33

Table 6 presents the results provided by the evaluators for the test cases 1 to 4 and with respect to the overall usefulness of BiOSS. These results show that the evaluators were very positive about the usefulness of the system. They have considered that the system is "Helpful" or "Very helpful". With respect to the output type, they think that the single output is more useful to solve real problems than the combined one.

Table 7. Evaluation results for the test case 5 (Synthetic Ontology).

Answer	No. of answers
The position of the Synthetic Ontology is adequate	5
The Synthetic Ontology should be ranked higher	1
The Synthetic Ontology should be ranked lower	0
Don't know	1
No answer	9

With respect to the test case 5 (see Table 7), the experts have considered that the result provided by the system (the Synthetic Ontology is ranked at the 3rd position) is adequate. In general, the experts have agreed on that the Synthetic Ontology should not be located at the first position because it is an ontology unknown by the biomedical community. However, they have considered that it should not be located at the last positions either, because it covers the 100% of the keywords provided as input.

Table 8. Other aspects of interest with respect to the evaluation of BiOSS.

Feature	Value
Average time used by the evaluators to rank ontologies without BiOSS (one test case)	Single output: 7 min. Combined output: 14 min.
Average time that the evaluators think they would require to provide an ideal ontology ranking without BiOSS (one test case)	Single output: 52 min. Combined output: 94 min.
Average time to complete the questionnaire (only the evaluators who filled out all sections were considered)	59 min.
Main resources and tools used by the evaluators to solve the tasks proposed in the questionnaire	BioPortal, OBO Foundry, Ontology Lookup Service, MeSH Browser, NCIt Browser, OntoBee, Google

Table 8 shows other aspects of interest regarding the evaluation process, such as the time used by the evaluators to rank the ontologies by hand, the time that the evaluators think it would be required to make an ideal ontology selection, etc.

Table 9. Main strengths of BiOSS according to the evaluators.

Strengths of BiOSS	No. evaluators
It is fast and precise	12
It is useful in fields where the user is not an expert and in fields that the user knows well but there exist a lot of ontologies (e.g. anatomy)	7
It is useful to confirm the expert's opinion	5
It is easy to use and the results are easy to understand	5
It shows the terms covered and not covered by the ontologies	4

Table 10. Main weaknesses of BiOSS according to the evaluators.

Weaknesses of BiOSS / Suggestions	No. evaluators
It should use BioPortal as the ontology repository	7
The level of formality of each source should be taken into account (e.g. ontology, controlled vocabulary, etc.)	4
It would be useful for the user to have links from each term to its definition and to related terms in the ontology	4
The URI for each ontology should be provided	3
Sometimes it provides erroneous results	2
The results contain some mistakes	2

Table 9 and Table 10 summarize the comments given by the evaluators on the main strengths and weaknesses of BiOSS. According to them, the main advantage of BiOSS is that it allows to

achieve ontology selection in a much faster and accurate manner than if done manually. As seen in Table 8, the average time estimated by the evaluators to carry out a "high-quality" ontology selection by hand is 52 minutes for a single test case and the single output. This time is much higher than the average execution time provided by the system (see Table 4). Experts also highlighted the great utility of the system in areas that are unfamiliar for them, and also in areas that they know but for which there are many different ontologies among which is difficult to choose the most appropriate.

As some recommendations to improve BiOSS, the experts have suggested to use the ontologies in BioPortal as the ontology repository. This improvement is proposed as a future work. Some of them have also stated that the level of formality of the knowledge represented by each source should be taken into account. They also consider that it would be very useful to link the results provided by the system to an ontology browser that would allow to explore the hierarchy of classes of each ontology and to display dynamically some useful information like relations, class properties and definitions.

6   CONCLUSIONS AND FUTURE WORK

Biomedical ontology selection is a difficult and time consuming task. The great number, size and complexity of biomedical ontologies make it really hard to choose the ontologies more adequate for a given domain or problem. This constitutes an important barrier to biomedical knowledge reusing. As a consequence, in the biomedical field it is becoming increasingly important to develop automatic methods and tools to select the best ontologies to describe particular datasets or to be used as the basis of specific applications (e.g. [37-43]).

A considerable amount of research has been done in the fields of ontology selection, search and evaluation. However, existing work consists mainly in theoretical approaches and methods. Currently, there is an important need for biomedical ontology selection tools and services that can be used to solve real-world problems.

The main contribution of this research is the release of BiOSS (Biomedical Ontology Selection System), a novel system for selecting biomedical ontologies. BiOSS is openly available at

http://bioss.ontologyselection.com. This article has described the system's functionalities as well as its architecture and workflow. In addition, an experimental evaluation of BiOSS and a comparison with other methods and tools in biomedicine have been presented.

According to our evaluation and revision of existing work, BiOSS can be considered one of the most complete biomedical ontology selection systems. Starting from a set of keywords, BiOSS processes them, extracts their meaning and achieves ontology selection without requiring user intervention. BiOSS is the only biomedical selection system that is able to evaluate ontology popularity and also the only one that has been designed to provide ontology combinations as output. It is available both as a web tool and a web service.

As the most immediate future work, BiOSS will be adapted to work with the BioPortal ontology repository, which is continuously maintained and updated, instead of using its own repository. The method to measure ontology popularity will also be improved. Popularity variation over time should be taken into consideration and some ontologies should be used even if they are not popular either because they are the reference standards or are required by regulation. In addition, we will review the manner of measuring ontology popularity in each Web 2.0 resource because according to the current method, there are situations in which an ontology could be considered popular even if it has negative references (e.g. if there are many tweets that provide negative comments about the ontology). We will also analyze the impact of taking into account other relevant ontology features, such as the frequency of update, the date of the last update and the level of formality. Finally, it is important to consider that some input terms may be more representative of the domain than others but BiOSS considers that all initial terms have the same importance. We will improve the user interface so that the user can assign different importance (weight) to different input terms.

ACKNOWLEDGEMENT

REFERENCES

[1] Y.A. Lussier, O. Bodenreider, Clinical Ontologies for Discovery Applications, in: B.C.J. O, C. K-H (Eds.) Semantic Web: Revolutionizing knowledge discovery in the life sciences, Springer, New York, 2007, pp. 101-119.

[2] M.A. Musen, N.F. Noy, N.H. Shah, P.L. Whetzel, C.G. Chute, M.-A. Story, B. Smith, The National Center for Biomedical Ontology, Journal of the American Medical Informatics Association, 19 (2012) 190-195.

[3] N.F. Noy, N.H. Shah, P.L. Whetzel, B. Dai, M. Dorf, N. Griffith, C. Jonquet, D.L. Rubin, M.A. Storey, C.G. Chute, BioPortal: Ontologies and Integrated Data Resources at the Click of a Mouse, Nucleic Acids Research, 37 (2009) W170-W173.

[4] J. Hartmann, R. Palma, A. Gómez-Pérez, Ontology Repositories, in: S. Staab, R. Studer (Eds.) Handbook on Ontologies, Springer, 2009, pp. 551-571.

[5] E. Simperl, Reusing Ontologies on the Semantic Web: A Feasibility Study, Data & Knowledge Engineering, 68 (2009) 905-925.

[6] D. Lonsdale, D.W. Embley, Y. Ding, L. Xu, M. Hepp, Reusing Ontologies and Language Components for Ontology Generation, Data & Knowledge Engineering, 69 (2010) 318-330.

[7] M. d'Aquin, N.F. Noy, Where to Publish and Find Ontologies? A Survey of Ontology Libraries, Web Semantics: Science, Services and Agents on the World Wide Web, 11 (2012) 96-111.

[8] M. Sabou, V. Lopez, E. Motta, V. Uren, Ontology Selection: Ontology Evaluation on the Real Semantic Web, in: Evaluation of Ontologies on the Web Workshop, held in conjunction with WWW'2006, Citeseer, Edinburgh, Scotland, 2006.

[9] A. Gómez-Pérez, Some Ideas and Examples to Evaluate Ontologies, Technical Report KSL-94-65. Knowledge Systems Laboratory. Stanford University. Also in Proceedings of the 11th Conference on Artificial Intelligence for Applications (CAIA94), (1994).

[10] A. Gómez-Pérez, From Knowledge Based Systems to Knowledge Sharing Technology: Evaluation and Assessment, Technical Report KSL 94-73. Knowledge Systems Laboratory, Stanford University, CA, (1994).

[11] T.R. Gruber, Toward Principles for the Design of Ontologies used for Knowledge Sharing, in: International Workshop on Formal Ontology. Available as Stanford Knowledge Systems Laboratory Report KSL-93-04, 1993.

[12] T. Berners-Lee, J. Hendler, O. Lassila, The Semantic Web, Scientific American, 284 (2001) 34-43.

[13] M. Sabou, V. Lopez, E. Motta, Ontology Selection for the Real Semantic Web: How to Cover the Queen's Birthday Dinner?, in: 15th International Conference on Knowledge Engineering and Knowledge Management. Managing Knowledge in a World of Networks. Lecture Notes in Artificial Intelligence, Springer, 2006, pp. 96-111.

[14] I. Cantador, M. Fernández, P. Castells, Improving Ontology Recommendation and Reuse in WebCORE by Collaborative Assessments, in: Workshop on Social and Collaborative Construction of Structured Knowledge at the 16th International World Wide Web Conference (WWW 2007), Citeseer, Banff, Canada, 2007.

[15] J. Brank, M. Grobelnik, D. Mladenic, A Survey of Ontology Evaluation Techniques, in: Conference on Data Mining and Data Warehouses (SiKDD 2005), Citeseer, Ljubljana, Slovenia, 2005.

[16] C. Jonquet, M.A. Musen, N.H. Shah, Building a Biomedical Ontology Recommender Web Service, Journal of Biomedical Semantics, 1 (2010) S1.

[17] J. Li, P. Shi, C. Hu, SOSEM: A Self-recommending Ontology Search Extension Mechanism Based on User Behavior, Journal of Software, 8 (2013) 1770-1777.

[18] N.F. Noy, P.R. Alexander, R. Harpaz, P.L. Whetzel, R.W. Fergerson, M.A. Musen, Getting Lucky in Ontology Search: A Data-Driven Evaluation Framework for Ontology Ranking, in: 12th International Semantic Web Conference (ISWC 2013).

[19] H. Tan, P. Lambrix, Selecting an Ontology for Biomedical Text Mining, in: Human Language Technology Conference, BioNLP Workshop Association for Computational Linguistics 2009, 2009, pp. 55-62.

[20] G. Maiga, A Flexible Biomedical Ontology Selection Tool, in: J.M. Kizza, K. Lynch, R. Nath, J. Aisbett, P. Vir (Eds.) Strengthening the Role of ICT in Development, Fountain Publishers, 2009, pp. 171-189.

[21] H. Alani, N.F. Noy, N. Shah, N. Shadbolt, M.A. Musen, Searching Ontologies Based on Content: Experiments in the Biomedical Domain, in: 4th International Conference on Knowledge Capture, K-Cap'07, ACM Press, 2007, pp. 55-62.

[22] O. Bodenreider, The Unified Medical Language System (UMLS): integrating biomedical terminology, Nucleic Acids Research, 32 (2004) 267-270.

[23] H. Alani, C. Brewster, Ontology ranking based on the analysis of concept structures, in: 3rd International Conference on Knowledge Capture (K-Cap), Banff, Canada, 2005, pp. 51-58.

[24] P. Buitelaar, T. Eigner, T. Declerck, OntoSelect: A Dynamic Ontology Library with Support for Ontology Selection, in: Demo Session at the International Semantic Web Conference, Citeseer, Hiroshima, Japan, 2004.

[25] C. Patel, K. Supekar, Y. Lee, E. Park, OntoKhoj: A Semantic Web Portal for Ontology Searching, Ranking and Classification, in: Workshop On Web Information And Data Management, ACM, 2003, pp. 58-61.

[26] L. Ding, R. Pan, T. Finin, A. Joshi, Y. Peng, P. Kolari, Finding and Ranking Knowledge on the Semantic Web, in: Y. Gil, E. Motta, V.R. Benjamins, M.A. Musen (Eds.) 4th International Semantic Web Conference, Springer, 2005, pp. 156-170.

[27] M. Martínez-Romero, J.M. Vázquez-Naya, C.R. Munteanu, J. Pereira, A. Pazos, An approach for the automatic recommendation of ontologies using collaborative knowledge, Lecture Notes in Artificial Intelligence, 6277 (2010) 74-81.

[28] M. Martínez-Romero, J.M. Vázquez-Naya, J. Pereira, A. Pazos, A Multi-criteria Approach for Automatic Ontology Recommendation using Collective Knowledge., in: Recommender Systems for the Social Web, Springer-Verlag, 2012, pp. 89-104.

[29] T.C. Rindflesch, A.R. Aronson, Ambiguity Resolution while Mapping Free Text to the UMLS Metathesaurus, in: The American Medical Informatics Society Annual Symposium on Computer Applications in Medical Care, American Medical Informatics Association, 1994, pp. 240-244.

[30] A. Jimeno-Yepes, A. Aronson, Knowledge-based biomedical word sense disambiguation: comparison of approaches, BMC bioinformatics, 11 (2010) 569.

[31] M. Martínez-Romero, Una Aproximación Multi-criterio para la Selección Automática de Ontologías Biomédicas Utilizando Conocimiento Colectivo, in: Department of Information and Communication Technologies, PhD Thesis, University of A Coruña, 122-133, 2011.

[32] B. Smith, M. Ashburner, C. Rosse, J. Bard, W. Bug, W. Ceusters, L.J. Goldberg, K. Eilbeck, A. Ireland, C.J. Mungall, The OBO Foundry: Coordinated Evolution of Ontologies to Support Biomedical Data Integration, Nature Biotechnology, 25 (2007) 1251-1255.

[33] M. Jones, H. Alani, Content-based Ontology Ranking, in: 9th International Protégé Conference, Stanford, CA, 2006.

[34] R. Studer, V.R. Benjamins, D. Fensel, Knowledge Engineering: Principles and Methods, IEEE Transactions on Data & Knowledge Engineering, 25 (1998) 161-197.

[35] H. López-Fernández, M. Reboiro-Jato, D. Glez-Peña, F. Aparicio, D. Gachet, M. Buenaga, F. Fdez-Riverola, BioAnnote: A Software Platform for Annotating Biomedical Documents with Application in Medical Learning Environments, Computer Methods and Programs in Biomedicine, 111 (2013) 139-147.

[36] M. Martínez-Romero, J.M. Vázquez-Naya, J. Pereira, M. Pereira, A. Pazos, G. Baños, The iOSC3 system: using ontologies and SWRL rules for intelligent supervision and care of patients with acute cardiac disorders, Computational and Mathematical Methods in Medicine 2013 (2013) 1-13.

[37] A.D. Torshizi, M.H.F. Zarandi, G.D. Torshizi, K. Eghbali, A Hybrid Fuzzy-Ontology Based Intelligent System to Determine Level of Severity and Treatment Recommendation for Benign Prostatic Hyperplasia, Computer Methods and Programs in Biomedicine, 113 (2013) 301-313.

[38] A. Gómez-Pérez, M. Martínez-Romero, A. Rodríguez-González, G. Vázquez, J.M. Vázquez-Naya, Ontologies in Medicinal Chemistry: Current Status and Future Challenges, Current Topics in Medicinal Chemistry, (2013).

[39] D. Isern, D. Sánchez, A. Moreno, Ontology-driven Execution of Clinical Guidelines, Computer Methods and Programs in Biomedicine, 107 (2012) 122-139.

[40] P. De Potter, H. Cools, K. Depraetere, G. Mels, P. Debevere, J. De Roo, C. Huszka, D. Colaert, E. Mannens, R. Van de Walle, Semantic Patient Information Aggregation and Medicinal Decision Support, Computer Methods and Programs in Biomedicine, 108 (2012) 724-735.

[41] C.M. Machado, D. Rebholz-Schuhmann, A.T. Freitas, F.M. Couto, The semantic web in translational medicine: current applications and future directions, Briefings in Bioinformatics, (2013) bbt079.

[42] J. Puentes, J. Montagner, L. Lecornu, J.-M. Cauvin, Information Quality Measurement of Medical Encoding Support Based on Usability, Computer Methods and Programs in Biomedicine, 112 (2013) 329-342.

[43] Y. Liu, A. Coulet, P. LePendu, N.H. Shah, Using Ontology-based Anotation to Profile Disease Research, Journal of the American Medical Informatics Association, 19 (2012) e177-e186.