



## Open Research Online

### Citation

Sabou, Marta and Pan, Jeff (2005). Towards improving web service repositories through semantic web techniques. In: Workshop on Semantic Web Enabled Software Engineering (SWESE) at ISWC'2005, 7 Nov 2005, ISWC.

### URL

<https://oro.open.ac.uk/3023/>

### License

None Specified

### Policy

This document has been downloaded from Open Research Online, The Open University's repository of research publications. This version is being made available in accordance with Open Research Online policies available from [Open Research Online \(ORO\) Policies](#)

### Versions

If this document is identified as the Author Accepted Manuscript it is the version after peer review but before type setting, copy editing or publisher branding

# Towards Improving Web Service Repositories through Semantic Web Techniques

Marta Sabou<sup>1</sup> and Jeff Pan<sup>2</sup>

<sup>1</sup> Vrije Universiteit, Amsterdam, The Netherlands

<sup>2</sup> University of Manchester

**Abstract.** The success of the Web services technology has brought topics as software reuse and discovery once again on the agenda of software engineers. While there are several efforts towards automating Web service discovery and composition, many developers still search for services via online Web service repositories and then combine them manually. However, from our analysis of these repositories, it yields that, unlike traditional software libraries, they rely on little metadata to support service discovery. We believe that the major cause is the difficulty of automatically deriving metadata that would describe rapidly changing Web service collections. In this paper, we discuss the major shortcomings of state of the art Web service repositories and, as a solution, we report on ongoing work and ideas on how to use techniques developed in the context of the Semantic Web (ontology learning, mapping, metadata based presentation) to improve the current situation.

## 1 Introduction

Web services technology relies on a stack of XML based protocols (such as WSDL<sup>3</sup>, SOAP<sup>4</sup>) that allow a uniform access to softwares running on different platforms and (even) implemented in different programming languages. The increased interoperability between heterogeneous software components allows their reuse and composition, thus leading to the high success of the Web services technology.

A prerequisite to reusing and composing Web services is the ability to find the right service(s). However, Web service discovery is becoming problematic with the increased number of Web services to several hundreds (e.g., there are already over 1000 Web services in bioinformatics [9]). Current research efforts investigate the possibility to automate Web service tasks (such as discovery and composition) by augmenting services with their formal semantic description<sup>5</sup> [11], [14], [20]. While this advanced technology is under development, the state of the art solution for finding Web services on the Web is inspecting online repositories of such services.

---

<sup>3</sup> Web Service Description Language, <http://www.w3.org/2002/ws/desc/>

<sup>4</sup> Simple Object Access Protocol, <http://www.w3.org/TR/soap/>

<sup>5</sup> See <http://www.daml.org/services/> for a number of initiatives.

In this paper we describe the major problematic aspects of online Web service repositories and show how Semantic Web related techniques could be used to enhance the situation. In particular, we investigate two research questions:

**Which are the problematic aspects of Web service repositories?** To answer this question we perform a survey of existing online Web service repositories and conclude on some of their major drawbacks (Section 2). We find that metadata acquisition and its meaningful presentation are two major problems that underly all negative aspects of these repositories.

**Which Semantic Web techniques could be used and how?** Metadata acquisition and its presentation are core issues for the Semantic Web. In this paper, we discuss how methods used in ontology learning (Section 3), ontology mapping (Section 4) and metadata based presentation (Section 5) could be useful to enhance the Web service repositories. The goal of this paper is to investigate the potential of Semantic Web related methods for supporting Web service discovery and to offer this overview as material for discussion during the workshop. We do not aim at presenting finalized research but rather to put forward our vision on the synergy between Semantic Web and (a particular aspect of) software engineering. Although some of our suggestions are only at an idea stage, most of them already have some support from previous experiments.

We summarize and point out future work in Section 6.

## 2 Web Service Repositories: State of the Art

In this section we tackle the first research question. For this (1) we summarize some major lessons learned from research on software libraries, (2) we perform an overview of online Web service repositories and (3) conclude on the major limitations of these repositories in comparison with software library standards.

### 2.1 Lessons Learned from Software Libraries

Storage and retrieval methods for software assets have been studied for almost three decades. A major survey of software reuse libraries concludes that, even if many sophisticated approaches exist to build and exploit such libraries, “*the practice is characterized by the use of ad-hoc, low-tech methods*” [12]. The practically viable approaches offer a good ratio between ease (and low cost) of implementation on one hand and a reasonable performance coupled with ease of use on the other.

From the six major types of approaches discussed by the survey, the *Information Retrieval* and *Descriptive* methods are the most widely used. Information retrieval methods regard software assets (source code, comments, design artifacts) as documents and adapt indexing techniques to these collections. Descriptive methods classify software assets in terms of a list of (predefined) keywords.

The most well known descriptive method is that of *faceted classification*, introduced by Pietro-Diaz [16]. In his approach, the keywords that describe the assets are organized per (possibly orthogonal) facets, thus defining a multidimensional search space (where each facet corresponds to a dimension). The survey considers that descriptive methods provide a better performance than information retrieval methods (in terms of precision and recall) and are easier to use. One of their drawbacks is that the acquisition of the right keywords and the classification of the assets according to these keywords increases the cost of their implementation [12].

## 2.2 An Overview of Online Web Service Repositories

In this section we overview seven Web services repositories. For each of them we describe the facilities that they offer to access the available services and point out problematic aspects when it is the case.

**1. UDDI**<sup>6</sup> is a cross-industry effort driven by major platform and software providers to establish an industry standard business registry which aims at facilitating Universal Description, Discovery and Integration of businesses and services. Different vendors (Microsoft, IBM, SAP) offer interfaces to this large repository. Hereby we present our findings for UDDI IBM and Microsoft. **UDDI-Microsoft**<sup>7</sup> allows both searching and browsing facilities. Browsing can be done according to several categorization schemes describing industry sectors (three versions of the North American Industry Classification System - NAICS), product catalogs (three versions of the United Nations Standard Products and Services Code - UNSPSC), geographic information (microsoft-com:geoweb:2000, ubr-uddi-org:iso-ch:3166-2003) and a small Web service classification scheme. This scheme contains 19 terms denoting domains (e.g., Health, Weather) and functionality types (e.g., Search, Printing). The search functionality is rather limited as it only allows searching for services whose name start with a given string. **UDDI-IBM**<sup>8</sup> provides a form based search (both for businesses and services), on the name of the services and a locator in one of the a categorization schemes.

**2. Bindingpoint**<sup>9</sup> is a repository of XML Web services. This site offers both search and browse facilities. Searching for a keyword will return any Web service which contains that keyword as a substring of the strings denoting the Web services name and description. This lack of tokenization (i.e., full text matching) leads to undesired effects. For example, when searching for “date”, any services that contain words such as “validate” or “update” (which are clearly not related to dates) are returned.

There is some ambiguity involved also in the presentation of the results. All resulting services are shown as well as the categories where a web service

<sup>6</sup> <http://www.uddi.org>

<sup>7</sup> <http://uddi.microsoft.com/>

<sup>8</sup> <https://uddi.ibm.com/beta/find>

<sup>9</sup> <http://www.bindingpoint.com/>

matching the search criteria was found. One would expect that when accessing these categories only those services which match the search criteria would be shown. This would be similar to performing a compound query, e.g., searching for all the *Calendar* type services that mention *date*. However, this is not the case - a click on the categories in this search context reveals all their members.

Browsing the available services can be done via two classification schemes. The BindingPoint scheme amounts to eight top categories which are further specialized up to two levels. The Visual Studio scheme consists of 15 categories without further specialization. Note, however, that even if some of these categories have the same name there is a considerable mismatch in their content. For example, the *Calendar* category has three instances in one classification scheme and twenty in the other.

**3. .NET XML Web Services Repertory**<sup>10</sup> offers a simple keyword based search on UDDI data using BindingPoint technology on both service names and descriptions. The same unsolicited results are obtained as in BindingPoint.

**4. WebserviceX.NET**<sup>11</sup> is a Web service provider that currently offers about 70 services. These services are grouped in seven categories which form the basic browsing mechanism.

**5. Web Service List**<sup>12</sup> provides 17 categories for browsing the available services (estimated 200). These categories denote the domain of Web services (e.g., Multimedia, Healthcare, Business/Finance) or a certain functionality they provide (e.g., Conversion, Search/Finders, Calculators). Besides, Web services can be browsed alphabetically, which is of little help if you do not know what you search for. Further, the site offers a search facility which searches the name and description of Web services. Unlike the BindingPoint technology this search works on correct tokenizations.

**6. Xmethods**<sup>13</sup> is one of the largest Web service repositories containing already several hundred services<sup>14</sup>. However, this site provides only a long list of services. It has no support for browsing nor does it provide any search facilities.

**7. SalCentral**<sup>15</sup> is a Web service repository which aggregates services published in other repositories (a meta-repository). It offers both search and a faceted classification based browsing. Searching is only performed on WSDL method names and textual service descriptions. However, search does not keep account of naming conventions of the composed method names. By considering each name as a string and simply performing substring search leads to several problematic cases. For example, searching for “text” could retrieve “GetGeoIP-Context” and “GetExtendedRealQuote”.

This repository is the only one that attempts a multi-facet based browsing. Services are classified according to six facets: the name of the method, country,

<sup>10</sup> <http://www.xmlwebservices.cc/>

<sup>11</sup> <http://www.webservicex.net>

<sup>12</sup> <http://www.webservicelist.com/>

<sup>13</sup> <http://www.xmethods.com>

<sup>14</sup> 425 on 19.07.2005

<sup>15</sup> <http://www.salcentral.com>

toolkit, domain, hosting server, suffix. Browsing is only possible on a single facet once, one cannot impose filters by selecting values from different facets. Since the values of the last 4 facets can be determined automatically they do not present any anomalies. However, we have several observations related to the first two facets.

The “**by method name**” facet offers a list keywords that frequently occur in the names of the methods offered by Web services. These keywords are:

Accounts, Address, Airport, Audio, Bill, Category, City, Credit, Client, Country, Currency, Customer, Database, Date, Domain, Email, Fax, File, Flight, Historical, Invoice, Location, Message, News, Postcode, Quote, Shop, Search, Sms, State, Supplier, Tax, Time, Town, User, Validate, Weather, Zipcode

It is not clear how these terms were derived, whether they have been simply manually selected or their selection involved some automatic analysis of the available services. There are several flows in this categorization:

**Incorrect instances.** Each category denoted by a keyword contains a set of Web services characterized by that keyword. However, this instantiation of categories with Web services is often incorrect - any Web service is a member of a category if the denoting keyword is contained in the name of the service. For example, when browsing the “Date” category we find instance services such as “*validateEmailAddress*” or “*updateAccountInfo*” which clearly should not belong to this category. This is a direct consequence of the employed search algorithm.

**Incomplete keyword set.** Several keywords have only a few instances (e.g., four instances for *Flight*). Nevertheless, terms that appear more often are missing from the offered keyword list. For example, searching for “text” returns four pages of results (about eighty hits) and searching for “phone” returns about 40 hits. This fact suggests that there is a mismatch between the terms frequently used by the collected services and those that are offered for browsing.

**Lack of abstractions.** Finally, many of these keywords are interrelated in a way that would allow grouping them in more generic (abstract) classes and building a deeper hierarchy to support browsing.

The “**by country**” facet offers a list of countries. The membership of a Web service to a country category is deduced based on the country extension of the URL or by using the manually added location information available in UDDI. Note, however, that the country extension of the URL does not always reflect the activity range of the service.

### 2.3 Conclusions

Based on the previously presented overview we conclude that the situation of Web service repositories is similar to that of software reuse libraries depicted a decade ago [12]. In particular:

**Simple techniques are used.** We encountered three simple ways of accessing the content of Web service repositories (see Table 1 for an overview). First, **search** is performed on (various combinations of) the textual sources attached to the Web services, such as their names, descriptions or the names of the WSDL operations (this is similar to the information retrieval methods implemented for software libraries). We encountered one case of searching where matching is done at token level (*keyword search*) and four cases where matching is done at substring level (*substring search*). Note that, substring search leads to many hits that have no content relevance for the search key. **Browsing** based on different categorization schemes (corresponding to descriptive techniques used in software libraries) is extensively used in Web service repositories. There are two types of schemes employed. First, large industry standard thesauri such as UNSPCSC and NAICS are used. These schemes are often under-populated and it is not always obvious which path to take to find what one needs. On the other hand, light-weight Web service specific classification schemes are also used. Finally, one repository uses no metadata to support browsing but simply presents all the available services as a large **list**.

**Browsing relies on few and low quality metadata.** Current Web service classification schemes are light-weight. Unlike the industry standard schemes, they have only a few top categories (max. 20) which, in most cases, are not further specialized. These schemes contain information about a single facet (except the case of SalCentral). Besides their reduced size and scope, Web service schemes are also qualitatively poor. For example, there is a high level of ambiguity of their scope since their categories often correspond to different facets. Some describe domains of activity (e.g., Health, Multimedia) while others name functionality types (e.g., Search, Find). Further, there is a mismatch between the content provided by the existing services and that covered by the categories. As a result, many categories are over-populated with instances and there is a need to extend the set of categories with new terms as the underlying data set evolves. Finally, it is often unclear how the categories are populated. In some cases, two identical categories are populated completely differently from the same set of services.

**The metadata is not fully exploited for presentation.** We found that sites which possessed richer metadata did not fully exploit this semantics for presentation. In particular, one of the advantages of faceted classifications is that they allow browsing on multiple facets at the same time (similar to a multi keyword search). However, current repositories allow only inspecting one facet at a time.

Searching for Web services is a complex issue. Their domain of activity is just one of the many criteria that characterize them. Especially when searching for services with the goal to reuse them in other applications it is important to know the functionality they offer, the type of input they require, the type of output that they produce, the restrictions that may apply on them.

Repository	Search	Browse	List
1. UDDI	- substring search	- product catalogs - one facet classification	
2. Bindingpoint	- substring search	- one facet classification	
3. .NET XML Web Services Repertory	- substring search		
4. WebserviceX.NET		- one facet classification	
5. Web Service List	- keyword search	- one facet classification	
6. Xmethods			Yes
7. SaCentral	- substring search	- six facet classification	

Table 1. Overview of retrieval methods in Web service repositories.

Web service repositories seldom use multi-faceted metadata. We believe that the major cause for this is the cost of acquiring these metadata. In the next two sections we will investigate how techniques from the fields of ontology learning and mapping could be used to acquire and enhance metadata for Web services. In Section 5 we will overview a few techniques that allow an intuitive presentation of faceted metadata. We will also show that some metadata we derived from our preliminary experiments can be successfully used as a basis for intuitive visualisations.

### 3 Ontology Learning

Ontology learning deals with developing methods for (semi-)automatically deriving ontologies from unstructured, semi-structured and structured data sets [10]. The stringent need of acquiring ontologies imposed by the development of the Semantic Web lead to the development of a large variety of approaches to this problem and already several tools that implement diverse ontology learning algorithms [15]. In previous work we successfully experimented with adapting existing ontology learning methods to deriving ontologies from textual Web service descriptions [17, 18]. In this section we show how ontology learning methods can be used to evolve Web service classification schemes or to extend them with new facets. We rely on some preliminary experimental results to strengthen the viability of some of our ideas.

#### 3.1 Evolve Existing Classification Schemes

The fast development of the Web services technology leads to a rapid growth in the number of available services. As a result many Web service categorization schemes lag behind the actual content needs of a dynamically changing collection of Web services. We experienced this phenomena during our survey of Web service repositories.

Our previous work showed that it is difficult for a domain expert to identify the terms that best describe a given collection of services [17, 18]. The reason is

that human experts do not perform a meticulous investigation of all available descriptions but rather rely on their own view of the domain to define the best terms. Our experiments showed that ontology learning techniques can support domain experts to identify the most frequent terms used by the community. Our proposal is to use concept identification to extend existing classification schemes and thus ensure that they truly reflect the content of the underlying repository.

For example, we ran our ontology learning module on a collection of services extracted from SalCentral and identified a set of terms (see below) that would extend the set of existing keywords in the “*by name*” category. We tested the relevance of these terms for the collection of Web services by searching for them in the collection. Several terms covered tens of services, thus proving their relevance for the collection.

text, temperature, stock, status, chart, company, word, price,  
payment, article, distance, language,

find, convert, verify, simulate, play, create, store, check,  
track, translate, calculate, validate.

Besides terms for broadening the existing scheme, our experiments also provided terms that would specialize existing keywords (i.e., “deepen” the scheme). Figure 1 depicts a few examples of specialization hierarchies that we learned.

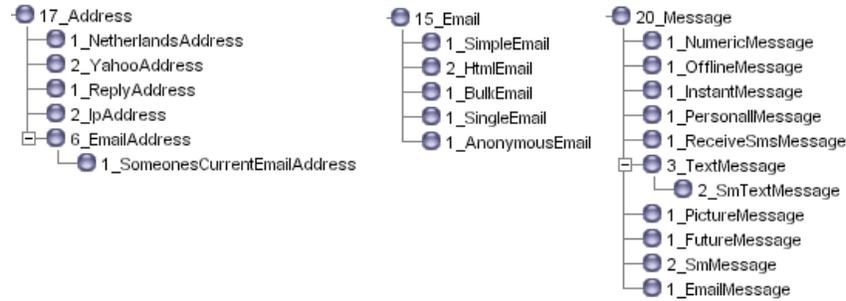


Fig. 1. Extracted specialization hierarchies.

### 3.2 Learn New Facets

In Subsection 2.3 we stressed the importance of using faceted classification schemes when describing services. We also observed that these faceted information is almost absent in current repositories probably due to its acquisition cost. During our experiments we found that the values for some of the basic

facets can be easily identified by using simple pattern matching techniques on the textual documentation of the services or inspecting their WSDL documentation. In this subsection we demonstrate our ideas about deducing operational features (input, output, functionality) and restrictions.

**Inputs, Outputs, Functionalities.** The type of input and output parameters as well as the action performed by a Web service are in many cases enough to identify the required service. While none of the analyzed repositories allow searching on these features, they can easily be identified with (semi-)automatic techniques.

First, the **textual descriptions attached to Web services** often contain this information. In fact, in our previous work we found that these texts exhibit very strong syntactic characteristics (they use a sublanguage [6]) and that this allows extracting the desired information by employing a few pattern based extraction rules. In particular, we observed that most of the noun phrases in these texts denote the parameters of the service while verbs indicate the functionality of the service. For example, in the following Web service descriptions the noun phrases *image*, *url address*, *hyperlink*, *web site*, *contact information*, *global address* denote the parameters of the service. The verbs *extract*, *validate* and *enhance* indicate the functionality of the service.

Extracts images from a given url address.

Extracts hyperlinks from a given web site.

Validate and enhance contact information for any global address.

Note that the above heuristics do not determine precisely which are the inputs and outputs of the service. For this, more refined rules can be defined. For example, “*given*” in front of a noun phrase indicates that it plays the role of an input. We are currently working on identifying such heuristics.

The second source of information for determining inputs, outputs and functionalities are **the WSDL files that describe Web services**. In particular the names of the methods and messages. From preliminary investigations it seems that WSDL files are often more accurate in providing this information than the textual descriptions. Our idea is that a combination of both sources should give the best results.

**Restrictions.** Besides operational features, such as inputs and outputs, other features can be important when choosing a service. In particular the geographic area where the service is active is an important consideration. Current repositories try to deduce this feature from the country extension of the URL where the service description is published. However, this seldom indicates the geographic region for which the service was built. For example, a Web service that “*validates*

and enhances contact information for any address in India” can be published at a .com address<sup>16</sup>. Conversely, a Web service whose URL contains a certain country identifier (e.g., France<sup>17</sup>) might perform a service that is independent of geographic constraints (e.g., in the case of the example service - cipher/decipher).

An alternative solution to determining geographic constraints for a service is to use Named Entity Recognition (NER) systems. Such systems automatically identify geographic entities, persons and organizations in free text. NER technology matured in the previous decades to reach performances of 80-90% Precision and Recall for a generic system (such as ANNIE) and 90-95% Precision and Recall for systems that are tuned to the needs of particular domains [3].

Search through all Swedish telephone subscribers.

Search UK Index.

This webservice return longitude, latitude and height from a given city. Only for France.

Lookup ATM Locations by Zip Code (US Only).

For example, for the Web service descriptions above, our experiments show that the ANNIE NER system recognizes *Swedish*, *UK*, *US*, *France* as references to the corresponding countries. We observed that, in some cases, the restriction is strengthened by the use of “only” in constructions such as “only for/in country” or “country only”. These constructs can be easily identified using a regular expression based rule mechanism.

### 3.3 Abstractions

The methods we presented so far identify important information about Web services. However, to be useful for browsing or even reasoning, these terms should be placed into subsumption hierarchies. There are several methods used to deduce subsumption relations. For example, in [2] four different techniques are combined to determine a subsumption hierarchy:

1. Hearst style lexico-syntactic patterns are matched against large corpora [7]. For example, the text snippet *carnivores such as lions, tigers* matches the pattern  $NP0 \text{ such as } NP1, NP2 \dots \Rightarrow isA(NP1, NP0), isA(NP2, NP0)$  and results in determining that lions and tigers are kinds of carnivores.
2. A similar approach is used to determine subsumption relations by taking advantage of the large amount of data offered by the World Wide Web. Given two terms, a set of Hearst like patterns are built up with them. The

<sup>16</sup> <http://ws.strikeiron.com/IndianAddressVerification?WSDL>

<sup>17</sup> <http://www.quisque.com/fr/chasses/crypto/cesar.asmx?WSDL>

occurrences of these patterns on the Web are counted and then normalized to determine the most likely relations.

3. WordNet is inquired for hypernymy information for the analyzed terms.
4. Vertical relations, as described in [19], are identified. This approach regards a term  $t1$  obtained by adding an extra modifier to a term  $t2$  as more specific than  $t2$ . For example, the term “*XML string*” is more specific than “*string*”.

Our ontology learning approach uses a vertical relations based algorithm to derive hierarchies of concepts that serve as parameters for the analyzed Web services. For example, by analyzing a collection of Web service from SalCentral we learned hierarchies as those depicted in Figure 1. The advantage of this very simple algorithm is that it performs well in terms of Precision (the majority of so identified subsumption relations are valid). The drawback is that it can only learn subsumptions indicated by compositionality (this results in a low Recall).

We also experimented with Hearst based patterns but these are very rare in the textual sources attached to Web services. For example, when analyzing around 450 descriptions, only 10 contained subsumption information identifiable with Hearst patterns. We will further explore the use of WordNet and the Web for hierarchy learning in this domain.

## 4 Ontology Reasoning and Mapping

In the previous section we presented a set of techniques to acquire metadata for characterizing Web services. In this section, assuming that we have ontologies, we briefly discuss how to make use of them in Web service repositories.

The domain ontology is crucial in tasks such as reasoning-based searches in Web service repositories. The basic idea is that the capabilities of Web services can be described by terms, in the domain ontology, that can be published, for example, by a UDDI registry. When clients search the repository, they can specify the capability of their desired Web services with a query term. The result of a search contains the Web services that are classified (by ontology reasoners such as Instance Store [1]) as instances of the query term. For example, [9] describes an application of reasoning-based search within bioinformatics.

However, there could be multiple domain ontologies related to a set of related Web service repositories. One of the challenges would be to harmonize these schemes by building a mapping between their terms. There is a variety of instance based mapping techniques that could be used for this purpose (see [13] for an overview). A well studied task for ontology mapping is to translate instances in the source ontology to instance to the target ontology, based on forward-chaining reasoning. It should be noted that, although well-studied, ontology mapping in general still presents many unsolved challenges, so that some care (and most likely some human involvement) would be needed in applying these techniques. Based on this task, a Web service repository can be extended by the increase of not only local registered Web services, but also ones in its federated Web service repositories.

## 5 Metadata Based Presentation

While important, the acquisition and reasoning of ontologies are just the first two steps of solving the problems of current Web service repositories. The intuitive presentation of ontology-based metadata is crucial to truly take advantage of its value. Faceted browsing and visual techniques are two frequently used ways to perform metadata based presentation.

### 5.1 Faceted Browsing

Several application domains have shown that (rich) faceted metadata provides a good basis for powering faceted browsing. For examples, faceted browsing interfaces were built to browse large image collections in the Flamenco project<sup>18</sup> [21] or to inspect museum item collections in the MuseumFinland project<sup>19</sup> [8]. This technology is reaching maturity as software vendors offer commercial products that automatically generate faceted interfaces from adequate metadata. A highly relevant example is the semantics-based Spectacle tool-suite offered by the Dutch company Aduna<sup>20</sup>.

The open source software repository, Sourceforge<sup>21</sup>, allows a faceted based browsing of the available applications. One can gradually narrow his search by imposing filters on the values of the available features. In the analyzed Web service repositories only SaCentral allows accessing different facets. However, there is no interaction between these facets — i.e., one cannot restrict the value of several facets at the same time.

Naturally, faceted browsing based portals can be easily built when the required metadata about Web services is in place.

### 5.2 Visualisation

Another way to present metadata is through visualisation techniques. Our previous work has shown that visualisation of faceted metadata can support several user tasks such as analysis, comparison and search [4]. We used Cluster Map [4], a visual technique developed by the Dutch company Aduna which is already integrated in several Semantic Web applications [5]. This technique visualizes *instances* of a set of *classes* according to their classification into these classes.

In this subsection we give an example of using Cluster Map to support the task of searching for Web services based on the automatically derived faceted metadata (using the previously described methods). Our current methods allow extracting two facets of the analyzed services: the types of their parameters and the functionality that they offer. These two facets are enough to answer queries that supply a functionality and a parameter type. For example, imagine that a user needs a service that finds addresses.

<sup>18</sup> <http://bailando.sims.berkeley.edu/flamenco.html>

<sup>19</sup> <http://museosuomi.cs.helsinki.fi>

<sup>20</sup> <http://aduna.biz/>

<sup>21</sup> <http://sourceforge.net/>

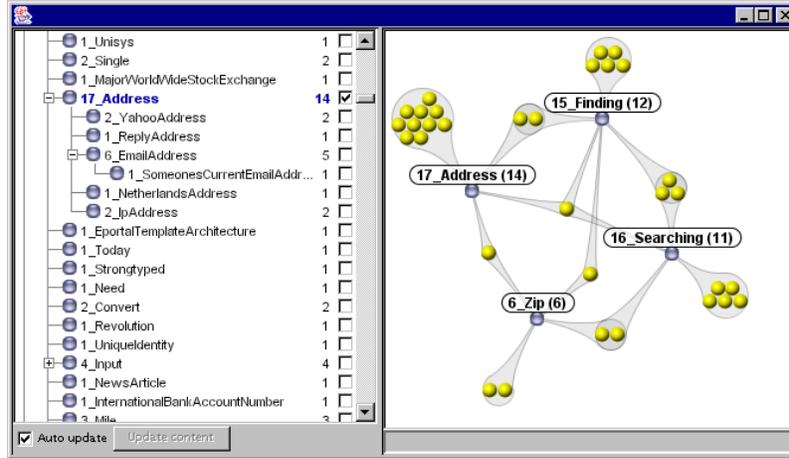


Fig. 2. An interface for visual search.

The Cluster Map technique is embedded in an interactive GUI as depicted in Figure 2. The left pane displays the hierarchy of terms. In this example, the hierarchy was automatically derived. The user of the interface can browse the hierarchy and select the terms that define his query. In the case of our query, the user might choose to see all services that offer *search* or *find* functionalities (from the *functionalities* facet). Also, he wants to see services that have parameters of type *address* and *zip* (these are values from the *parameters* facet). Note that, by displaying all the domain relevant terms, we offer support for formulating the users query in terms that are actually used within the service collection.

The selected terms are visualised in the right pane with their name and cardinality stated in a rounded rectangle. Balloon-shaped edges connect instances (small yellow spheres) to their most specific class(es). In this case the instances of a term are all the Web services that are described by that term. Instances with the same class membership are grouped in clusters (similar to Venn Diagrams). In our example, there are several clusters formed, one of them showing the intersection between *Address* and *Finding*. This cluster contains two Web services which have a parameter of type *Address* and perform the action of *Finding* - thus they represent the answer to our example query. The instances in a cluster can be accessed with a mouse click.

This visualisation allows the user to explore the service collection. For example, in the example scenario, he might be interested to see what other services provide find functionalities, or to inspect the one service that allows finding zip codes. Further, using the specialization hierarchy in the left pane he can refine the query, for example by choosing more specialized terms (in our case, he might actually be interested in email addresses).

## 6 Summary

In this paper we investigated the use of various Semantic Web related techniques to enhance current online Web service repositories.

Our overview of Web service repositories yielded that they rely on little and qualitatively poor metadata. As a consequence they offer only limited support for performing manual Web service discovery. Inspired by the lessons learned from traditional software libraries, we believe that the use of rich faceted metadata would be required. However, we are aware that acquiring such metadata, especially for describing Web service collections that are changing on a daily bases, is prohibitively expensive.

As a solution, we think that techniques developed for the Semantic Web which are concerned with metadata acquisition and presentation have a great potential in solving the current problems of Web service repositories. In particular, ontology learning techniques can be adapted for extending and keeping up to date the current service classification schemes. They can also be used to derive the information for several facets that describe services or to arrange the extracted terms in meaningful subsumption hierarchies. We already presented some encouraging results when using the ontology learning techniques. Ontology reasoning techniques can be used to provide semantic-based search of Web services. Ontology mapping techniques could be used to support federated Web service repositories. Finally, rich faceted metadata can be exploited for building intuitive browsing interfaces. We exemplified that even the light-weight metadata that we derived automatically can significantly enhance the search for Web services when coupled with visualisation techniques.

Encouraged by our results so far, we prepare to implement all these ideas in a prototype system that would collect available Web services, automatically extract metadata describing different facets of these services and then would use this metadata to build an intuitive search/browse interface. Furthermore, we believe that a minimal user intervention would be enough to “clean” the automatically derived metadata so that it can be used as a basis for (simple) reasoning tasks. This would bring the state of the art a step closer to the ultimate vision of Semantic Web services where semantics is added in a bottom-up fashion (learned from available sources) rather than being imposed in a top-down approach (requiring costly manual annotations).

## References

1. Sean Bechhofer, Ian Horrocks, and Daniele Turi. The OWL instance store: System description. In *Proc. of the 20th Int. Conf. on Automated Deduction (CADE-20)*, Lecture Notes in Artificial Intelligence. Springer, 2005. To appear.
2. P. Cimiano, A. Pivk, L. Schmidt-Thieme, and S. Staab. Learning Taxonomic Relations from Heterogeneous Evidence. In *Proceedings of the ECAI04 Workshop on Ontology Learning and Population*, Valencia, Spain, 2004.
3. H. Cunningham, D. Maynard, K. Bontcheva, and V. Tablan. GATE: A Framework and Graphical Development Environment for Robust NLP Tools and Applications.

- In *Proceedings of the 40th Anniversary Meeting of the Association for Computational Linguistics*, Philadelphia, July 2002.
4. C. Fluit, M. Sabou, and F. van Harmelen. Ontology-based Information Visualisation. In V. Geroimenko, editor, *Visualising the Semantic Web*. Springer, 2002.
  5. C. Fluit, M. Sabou, and F. van Harmelen. Ontology-based Information Visualization: Towards Semantic Web Applications. In V. Geroimenko, editor, *Visualising the Semantic Web, Second Edition*. Springer, 2005.
  6. R. Grishman and R. Kittredge, editors. *Analyzing Language in Restricted Domains: Sublanguage Description and Processing*. Lawrence Erlbaum Assoc., 1986.
  7. M.A. Hearst. Automatic Acquisition of Hyponyms in Large Text Corpora. In *Proceedings of the Fourteenth International Conference on Computational Linguistics*, 1992.
  8. E. Hyvonen, E. Makela, M. Salminen, A. Valo, K. Viljanen, S. Saarela, M. Junnila, and S. Kettula. MuseumFinland – Finnish Museums on the Semantic Web. *Journal of Web Semantics*, 2005. To appear.
  9. P.W. Lord, P. Alper, C. Wroe, and C.A. Goble. Feta: A Light-Weight Architecture for User Oriented Semantic Service Discovery. In A. Gomez-Perez and J. Euzenat, editors, *The Semantic Web: Research and Applications. Proceedings of the Second European Semantic Web Conference, ESWC'05*, volume 3532 of *LNCS*, pages 17–31, Heraklion, Crete, Greece, May 29 – June 1 2005. Springer-Verlag.
  10. A. Maedche and S. Staab. Ontology Learning for the Semantic Web. *IEEE Intelligent Systems*, 16(2):72–79, March/April 2001.
  11. S. McIlraith, T.C. Son, and H. Zeng. Semantic Web Services. *IEEE Intelligent Systems. Special Issue on the Semantic Web*, 16(2):46–53, March/April 2001.
  12. A. Mili, R. Mili, and R.T. Mittermeir. A survey of software reuse libraries. *Annals of Software Engineering*, 5:349 – 414, 1998.
  13. N.F. Noy. Semantic Integration: A Survey Of Ontology-Based Approaches. *ACM SIGMOD Record*, 33(4):65 – 70, December 2004.
  14. M. Paolucci, T. Kawamura, T. Payne, and K. Sycara. Importing the Semantic Web in UDDI. In *Proceedings of the Workshop of Web Services, E-Business, and the Semantic Web*, 2002.
  15. A. Gomez Perez and D. Manzano Mancho. A Survey of Ontology Learning Methods and Techniques. *OntoWeb Deliverable 1.5*, May 2003.
  16. R. Pietro-Diaz. Implementing Faceted Classification for Software Reuse. *Communications of the ACM. Special issue on software engineering.*, 34(5):88 – 97, 1991.
  17. M. Sabou. From Software APIs to Web Service Ontologies: a Semi-Automatic Extraction Method. In *Proceedings of the Third International Semantic Web Conference, ISWC*, Hiroshima, Japan, November 2004.
  18. M. Sabou, C. Wroe, C. Goble, and G. Mishne. Learning Domain Ontologies for Web Service Descriptions: an Experiment in Bioinformatics. In *Proceedings of the 14th International World Wide Web Conference*, Chiba, Japan, May 2005.
  19. P. Velardi, M. Missikoff, and P. Fabriani. Using Text Processing Techniques to Automatically enrich a Domain Ontology. In *Proceedings of the International Conference on Formal Ontology in Information Systems*, pages 270–284, 2001.
  20. M. Voskob. UDDI Spec TC V4 Requirement - Taxonomy support for semantics. OASIS, 2004.
  21. K.-P. Yee, K. Swearingen, K. Li, and M. Hearst. Faceted Metadata for Image Search and Browsing. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 401 – 408, Florida, USA, 2003. ACM Press.