Semi-automatic Semantic-Based Web Service Classification

Miguel Ángel Corella and Pablo Castells

Universidad Autónoma de Madrid, Escuela Politécnica Superior
Campus de Cantoblanco, 28049 Madrid, Spain
{miguel.corella, pablo.castells}@uam.es

Abstract. With the expectable growth of the number of Web services available on the WWW and service repositories, the need for mechanisms that enable the automatic organization and discovery of services becomes increasingly important. Service classification using standard or proprietary taxonomies is a common and simple facility in this context, complementarily to more sophisticated service management retrieval techniques. In this paper we propose a heuristic approach for the semi-automatic classification of Web services, based on a three-level matching procedure between services and classification categories, assuming a corpus of previously classified services is available. An experimental test of the proposed techniques is reported, showing positive results.

1 Introduction

Since the emergence of the semantic Web [3], many research efforts have been aiming to use semantics to endow Web services with a higher potential for automation. These efforts have given rise to the semantic Web services vision [18]. The basis of this trend is to add semantic information to current Web service descriptions (in WSDL [6] format) to enable their analysis and manipulation by software programs enacting further automation capabilities for such tasks as service selection, invocation, composition, or discovery [11] and other tasks related to Web services but not often mentioned as being target of semantic-based technologies. In this paper we focus on the classification of Web services, this is, the assignment of a class to a service, at publication time, indicating the domain (i.e. the business focus) to which a service belongs.

Today, UDDI [14] is the most widely accepted and used protocol for publishing and searching Web services on the Web. These actions are usually performed within UDDI registries, which can be defined as service repositories available (and easy accessed through a URL) on the Internet. In these registries, services can be classified using one or several service taxonomies (such as UNSPSC[1] – United Nations Standard Products and Service Code, NAICS[2] – North American Industry Classification System, or even user-defined taxonomies created in UDDI format, since UDDI specification v.2). Nevertheless, this classification has to be performed manually by a

---

[1] http://www.unspsc.org/
[2] http://www.census.gov/naics

human publisher. Due to the huge quantity of classes in standard service taxonomies like the ones already mentioned, the classification process is usually costly. Furthermore, taxonomies are subject to evolution, change or even replacement by new ones, making even heavier the maintenance effort load on repository administrators.

The purpose of the work presented here is to provide automatic mechanisms to help service publishers in the classification task at publication time. For this aim, we propose a heuristic-based classification system that compares a new service with the ones already classified in order to predict the appropriateness of the available classification categories for the new service, and produce a ranked list of candidate classes. Service classification is not only a must for manually browsing and managing service repositories, but can in fact be used for simple but efficient forms of service annotation and semi-automated discovery, complementary aid for automatic selection, etc.

The paper is organized as follows: Section 2 introduces some related work already presented in the domain of Web service classification and other related knowledge areas. Section 3 defines the problem of Web service classification in a more formal way and motivates why service semantics are needed in order to successfully solve it. The presentation and brief explanation of our classification heuristic is described in Section 4. Section 5 describes a framework where the classification method has been implemented and tested. Section 6 reports the experimental results obtained with our approach. Finally, Section 7 provides conclusions and outlines future work directions.

2   Related Work

The problem of the automatic classification of Web services has been addressed in prior work from two main approaches, that we may class as heuristic (e.g. [15]) and non-heuristic (e.g. [5] and [10]). In [10], two different strategies are proposed. The first one is based on using the information contained in non-semantic service descriptions to select a category in which the service fits best, by using Natural Language Processing, machine learning and text classification techniques. The second one consists of using the same information contained in service descriptions to dynamically create the categories in which the service should be classified, using clustering techniques. In both approaches the classification process is based on the extraction of relevant words from service descriptions, the construction of term vectors with those relevant words and the usage of classification mechanisms (e.g. Naïve Bayes) to perform the vector classification. This way, the service classification problem is solved by a text classification approach. From our point of view, this approach has two main drawbacks. First, the hypothesis of finding relevant and meaningful words in service descriptions is a very optimistic starting point. Next, doing clustering implies that the created categories do not have meaningful names, and the classification taxonomy changes over time. These two problems do not matter at publication time, but can be an issue if the classification information is intended to be used for service discovery, since users could neither select services by category (they do not have a name), nor get properly acquainted with the taxonomy, as its structure may change frequently.

The approach to classification proposed in [5] follows similar steps as those described in [10]. The main difference is that the method used to classify term vectors is based on Support Vector Machines. In addition, in this proposal, service publishers are provided with some extra information, more precisely, with a concept lattice extracted using Formal Concept Analysis over service descriptions. This extra information allows service publishers to know how the words used in their service descriptions contribute to the selection of a specific category, helping them to e.g. modify some words of their descriptions which may cause ambiguity in the classification process. As this approach also applies Natural Language Processing techniques on service descriptions, the same drawback as pointed out above can be found here, namely the assumption that meaningful textual information can be found in service descriptions (operation and parameter identifiers, comments, etc.) does often not hold.

In [15], a framework to semi-automate the semantic annotation of Web services (i.e. parameter description based on ontology concepts, semantic service classification, etc.) is presented. For classification, an algorithm to match Web service data types (in XML Schema[3]) and domain ontology concepts is defined based on schema matching. Assuming a 1-1 correspondence between domain ontologies and service categories, the classification is done by selecting the category that corresponds to the domain ontology that yields the highest similarity when compared to the service. The drawback here is that best practices in Web service definition prescribe document-based service descriptions, this is, service messages should consist of a unique part defined by a complex schema containing all the service parameters. With this form of description it is difficult to find similarities with domain ontology concepts as, usually, no single domain concept will contain the complete structure of service messages.

Another relevant area for our work is that of service matchmaking (see e.g. [12] and [16]). This research topic is related to Web service classification in that our approach to service classification computes similarity degrees between services in order to assign them a common category, and service matchmaking aims to find services that match a concrete capability description, e.g. in order to invoke matching services and/or compose them into more complex processes. The main difference between this research area and our addressed problem is that while in our view, service classification admits some degree of fuzziness in service matching, i.e. we consider a continuous similarity measure, work on service matchmaking typically does not, and is based on discrete matching levels (e.g. "no match", "partial match", "complete match").

3   The Service Classification Problem

As mentioned earlier, service classification is a common necessity to make service administration and retrieval manageable for human users. Moreover, it can serve as a complementary aid for automatic service discovery and selection techniques. Nevertheless, there are often usability problems involved in service categorization which cause difficulties for the creation, validation, and use of classifications in real-world environments. Such problems include:

---

[3] http://www.w3.org/TR/2001/REC-xmlschema-2-20010502

- Classification taxonomies can be extremely large, comprising thousands of categories (e.g. UNSPSC ~ 20,000 classes, NAICS ~ 2,300 classes).
- The number of services in a repository can grow quite large, making it virtually impossible for repository administrators to validate the information published along with a service.
- The placement of a service under a proper category requires a considerable amount of knowledge of the taxonomy, the application domain, etc., in order to make appropriate classification decisions. Few publishers or administrators have this knowledge with sufficient width and depth.

Our work aims at alleviating the administrator's work, and reducing the publication effort for service providers, by supplying them with a ranked set of likely appropriate categories when a new service has to be published in the repository. Our proposal approaches the classification problem as follows. Given a set of services already classified under a given taxonomy, and a new service description to be published, the unclassified service is compared with the classified ones, whereby a measure of the likelihood that the service should be assigned a certain category is computed.

WSDL descriptions provided by current technologies are not suitable for this purpose, as they only focus on the syntactic view of the services, which is not sufficient to support valid service classification criteria in practice.

3.1   The Need for Service Semantics

Consider this example: take two Web services, the first one defining currency conversion capabilities, and the second one describing a trip time calculator. A typical description of such functionalities in WSDL is provided in Appendix A at the end of this paper. Since WSDL service descriptions only include functional information (i.e. syntactic information about the service interface), the only available description elements to compare services are service operations, messages, and data types. As can be seen in the example, the currency converter service has one operation, involving:

- An input message containing an amount of money of type double, and two currency codes, of type string.
- An output message containing the converted amount (a double).

On the other hand, the trip time calculator service has also one operation, involving:

- An input message containing an average speed of type double, and two city names, of type string.
- An output message containing the trip time in minutes (a double).

From a conceptual point of view, these two services should yield a low similarity measure value when compared. However, since the WSDL interfaces are syntactically equivalent, their comparison would produce a very high result value.

- In conclusion, WSDL-based descriptions are not sufficient, and would often lead to inconsistent similarity values, and therefore, to service misclassification. Semantic Web service descriptions can solve this problem by providing means to describe service inputs and outputs from a conceptual point of view. A typical description

of the examples as semantic Web services is included in Appendix B. We use WSMO [17] in the example, but our classification approach (as it is based on abstract mathematical similarity formulas) is language agnostic, so it is compatible with other semantic Web service languages such as OWL-S [13], WSDL-S [1], or SWSO [2].

4   Classification Heuristic

As introduced in earlier sections, our heuristic approach consists of the comparison of unclassified services with classified ones. The heuristic is divided into three levels, corresponding to the comparison between different service elements involved in the classification procedure, as we explain next. We will omit here all the mathematical formalization details, which can be found in [7].

Service Category Level. Since services have to be assigned a category as a result of the classification procedure, the consideration of this level is quite obvious. It is needed in order to find evidence that a service should belong to a specific category. This is the highest level of granularity in the classification method, at which a final service-category matching degree is obtained, which is used to sort the ranked service category list proposed to service publishers. The proposed computation for this measure is defined by:

$$\mathrm{P}(s:c) \sim \sum_{A \subset \tau^{-1}(c)} (-1)^{|A|+1} \prod_{x \in A} \mathrm{sim}(s,x)$$

where P(s:c) is the evidence that the service s should belong to the category c formalized in a probabilistic way and estimated by the inclusion-exclusion principle [19] applied to a set of computed similarity values sim(s,x), between service s and all the services under category c. Thus, the predicted appropriateness of a category for a service increases with the similarity between the service and the classified services.

Service Description Level. The comparison between services is based on the assumption that services of the same category are likely to deal with similar concepts as inputs/outputs. Therefore, operation structures (i.e. conceptual roles and grouping of the data involved in the operations) are considered relevant for service-level comparisons. In fact, the similarity between two services is measured in terms of the similarity between service operation sets, and it is computed as the average of the best possible pairwise similarities obtained by an optimal pairing of the operations from the two sets. For this purpose, the similarity between two operations is computed as:

$$\mathrm{sim}(op,op') = \mathrm{sim}(I_{op}, I_{op'}) \cdot \mathrm{sim}(O_{op}, O_{op'})$$

where I_{op}, $I_{op'}$, O_{op} and $O_{op'}$ are the set of input and output parameters of the operations op and op' respectively. The similarity between two parameter sets is computed in turn as the average of the best possible pairwise similarities obtained by an optimal pairing of the parameters from the two sets.

Service Parameter Level. The comparison of service parameters in our approach is based on the similarity between the ontology concepts used to annotate them. A considerable body of research has addressed the problem of matching ontology concepts (e.g. [4], [8], [9]). In our current experiments, we have tested our own concept to concept similarity measure, specifically devised and tuned to our approach. The similarity between two ontology concepts t and t' is measured by:

$$\mathrm{sim}(t,t') = \left(1 - \frac{\alpha}{\mathrm{h}(\mathcal{T})} \cdot \frac{|d-d'|}{d+d'}\right) \cdot \frac{1}{\min(d,d')} \cdot \left(1 - \frac{\max(d,d')-1}{\mathrm{h}(\mathcal{T})}\right)$$

where $\mathrm{h}(\mathcal{T})$ is the total height of the concept hierarchy in the ontology, d and d' are the distances from t and t', respectively, to their lowest common ancestor, and $\alpha \in [0,1]$ is a parameter that ensures a minimum non-zero similarity value to soften the impact of this measure on the overall heuristic. In our tests, α was tuned empirically to 0.8.

5   Implementation

We have implemented the techniques described in the previous sections into a service classification framework, serving the double purpose of a) demonstrating a practical environment where users can classify and store their services in a repository, and b) setting up experiments to test and evaluate our approach .
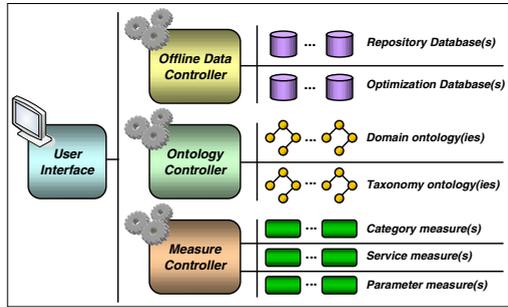
Fig. 1. Classification framework architecture showing the three main component types of the implementation: offline data controllers, ontology controllers and measure controllers

Figure 1 shows a high-level view of the main architectural components of this framework. These include:

- Offline Data Controllers: They are responsible for controlling the access to offline information stored in different databases. This offline information storage allows both the persistence of the service repository and the optimization of the execution time. For instance, since our techniques involve a combinatory comparison between services, parameters, etc., many similarity values are computed in advance and

stored in the database, The offline data controllers are compatible with any service repository, by implementing a bridge from the repository to our database.

- Ontology Controllers: They interface with the different ontologies that may be involved in the classification process. These include both domain ontologies containing the concepts used in semantic Web service descriptions, and service taxonomies, represented in an ontological format. Again the generality of these components allows plugging any ontology into the framework.
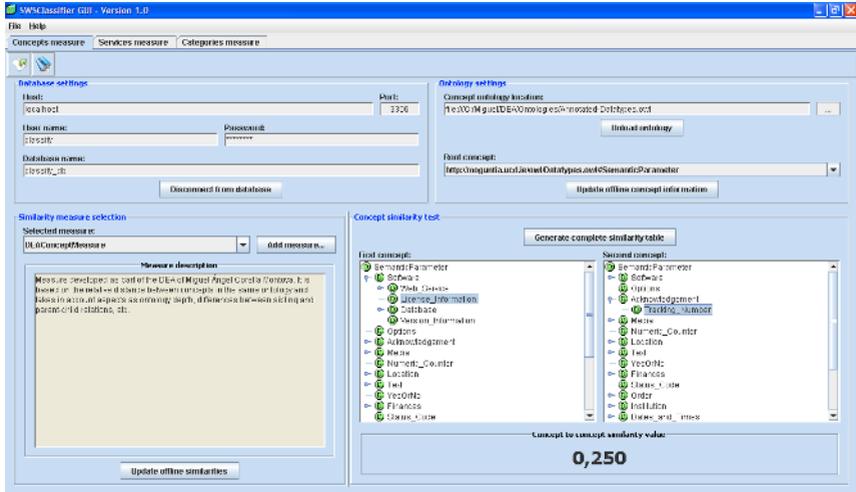

**Fig. 2.** Screenshot of the classification framework to use and test our approach

- Measure Controllers: They interface with the similarity measures involved in the heuristic. Once again, these abstraction components enable using any similarity measure at each level provided that it complies with the controller interface.

The above components are provided to achieve as much generality as possible. This way, our framework can be easily configured for the classification of services described with any domain ontology, into any taxonomy, based on any matching functions (at each comparison level), and using in any repository implementation standard (e.g. UDDI). This flexibility is also a valuable feature to facilitate the tests involved in our research. A screenshot of the framework user interface is shown in Figure 2.

6   Experiments and Evaluation

The approach proposed here was tested and evaluated in the implemented framework. The corpus used in the experiments is a repository containing 164 semantic Web service descriptions (in OWL-S), annotated using an ontology containing over 400 atomic concepts (i.e. no complex concept definitions, as e.g. OWL or WSML support, were used yet in our experiments). The services were manually classified using a taxonomy containing 23 different service categories. The service repository was

essentially the one used in several other studies [5, 10, 15], and available in A. Heβ's Web page[4]. We have reused the domain ontology and the taxonomy included in this corpus, but have mapped the service descriptions to WSMO in order to reuse some tools available or developed in our research group (e.g. an input/output concept extractor or a translator to Racer[5] logical axioms enabling reasoning in our future work).
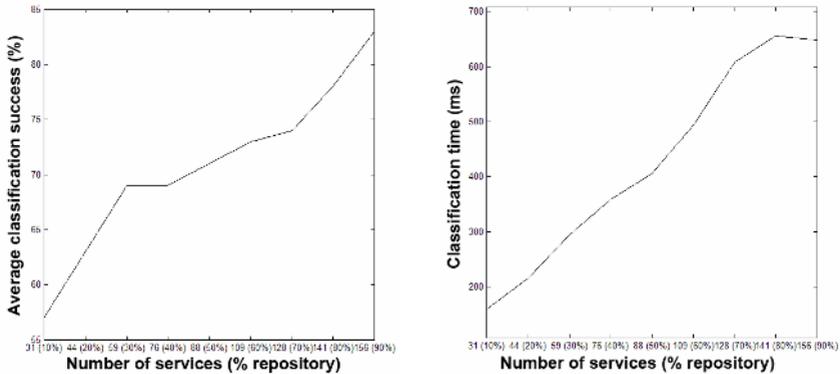
Fig. 3. Average classification success rate (i.e. correct class is at top of the ranking) vs. the services used as evidence (left), and average service classification time vs. the number of services in the repository (right). The test was performed with a repository of 164 services.

Using this corpus, we have conducted several effectiveness and performance tests by using a part of the repository as evidence (i.e. classifications taken as a given), and the rest for testing (i.e. services to be classified). The experiments were run for different ratios of evidence vs. test corpus sizes, in order to observe the evolution of the performance measures with respect to the amount of training data. The results obtained in the experiments are shown in Figure 3. The tested features are:

- Classification success rate: In this test we measured the average percentage of correct classifications with respect to the percentage of the repository used as evidence and test. A correct classification is one that ranks first the right (manually assigned) category of a service in the ranking of possible categories. This provides an estimation of the average probability of offering the correct category as first option. It can be seen that the success rate is about 83% when almost all the available corpus of 164 (but eight) services is used as evidence. In the cases where the method failed, the correct class was usually second or third in the ranking.
- Time performance: Although this is not a critical issue for classification, since this process can be performed offline at service publication time, our framework aims to be as efficient as possible. Thus we have measured the average classification time with respect to the size of the repository, which shows linear growth up to less than 7 s., on an Intel® Pentium® M, 1.73 GHz and 1 GB of RAM available.

---

[4] http://www.few.vu.nl/~andreas/projects/annotator/owl-ds.html
[5] http://www.racer-systems.com

7   Conclusions and Future Work

We have presented an efficient service classification approach based on conceptual service descriptions, that can be used to assist publishers, consumers and repository administrators in manual service categorization and retrieval tasks. As a continuation of the work presented here, we plan to investigate the potential of the proposed classification capabilities to enhance automatic service retrieval mechanisms. The heuristics have been tested on a corpus use in prior research by different authors, showing positive results. The generality of the implemented framework allows the easy integration and testing of different similarity measures at the three proposed granularity levels.

Besides the combination of our approach with service discovery techniques, as future work we envisage the extension of the algorithms to deal with more complex ontology-based descriptions of concepts, service capabilities (by Boolean expressions), etc. This involves a generalization of our basic matching functions, which could benefit from available ontology-oriented reasoners, and could link to ongoing related research in service matchmaking. Further experimentation and testing of our approach is foreseen as well, such as the comparison with other existing techniques, similarity measures, as well as the combination of several measures into improved ones, tests on larger repositories, performance tests with respect to further corpus features (such as taxonomy and ontology size, service disparity), etc.

Acknowledgements

References

1. Akkiraju, R., Farrel, J., Miller, J., Nagarajan, M., Schmidt, M., Sheth, A., Verma, K: Web Service Semantics – WSDL-S, Technical Note, Version 1.0, 2005.
2. Battle, S., Bernstein, A., Booley, H., Grosof, B., Gruninger, M., Hull, R., Kifer, M., Martin, D., McIlraith, S., McGuiness, D., Su, J., Tabet, S.: Semantic Web Service Ontology (SWSO), Version 1.0, 2005.
3. Berners-Lee, T., Hendler, J., Lassila, O.: The semantic web. Scientific American, 2001.
4. Bernstein, A., Kaufmann, E., Bürki, C., Klein, M.: How similar is it? Towards personalized similarity measures in ontologies. In the 7[th] Internationale Tagung Wirtschaftsinformaitk. Bamberg, Germany, 2005, pp. 1347-1366.
5. Bruno, M., Canfora, G., Di Penta, M., Scognamiglio, R.: An approach to support web service classification and annotation. In Proceedings of the IEEE International Conference on e-Technology, e-Commmerce and e-Services (EEE 2005), Hong Kong 2005.
6. Christiensen, E. et al: Web Service Description Language (WSDL), v1.1.

7. Corella, M. A., Castells, P.: A Heuristic Approach to Semantic Web Services Classification. In Proceedings of the 10th International Conference on Knowledge-Based & Intelligent Information & Engineering Systems (KES 2006), Bournemouth, UK, 2006.

8. Culmore, R., Rossi, G., Merelli, E.: An ontology similarity algorithm for BioAgent. In NETTAB 02 Agents in Bioinformatics. Bologna, Italy, 2002.

9. Ehrig, M., Haase, P., Stojanovic, N.: Similarity for ontologies – a comprehensive framework. Workshop on Enterprise Modelling and Ontology at PAKM 2004. Austria, 2004.

10. Heβ, A., Kushmerick, N.: Automatically attaching semantic metadata to Web Services. In Workshop on Information Integration on the Web (IIWeb2003), Acapulco, Mexico, 2003.

11. Keller, U., Lara, R., Lausen, H., Polleres, A., Fensel, D.: Automatic Location of Services. In 2nd European Semantic Web Conference (ESWC 2005). LNCS Vol. 3532 pp. 1-16.

12. Li, L., Horrocks, I.: A software framework for matchmaking based on semantic web technology. In the International Journal of Electronic Commerce, 8(4):39 – 60. 2004.

13. Maritn, D., Burnstein, M., Hobbs, J., Lassila, O., McDermott, D., McIlraith, S., Narayanan, S., Paolucci, M., Parsia, B. et al: OWL-S: Semantic markup for web services, v1.1, 2004.

14. OASIS: UDDI: The UDDI technical white paper, 2004.

15. Oldham, N., Thomas, C., Sheth, A., Verma, K.: METEOR-S Web Service Annotation Framework with Machine Learning Classification. In Proc. of the 1st Int. Workshop on Semantic Web Services and Web Process Composition (SWSWPC'04), California, July 2004.

16. Paolucci, M., Kawamura, T., Payne, T., Sycara, K.: Semantic Matching of Web Service Capabilities. In Proceedings of the First International Semantic Web Conference, 2002.

17. Roman, D., Lausen, H., Keller, U., de Brujin, J., Bussler, C., Domingue, J., Fensel, D., Hepp, M., Kifer, M., König-Ries, B., Kopecky, J., Lara, R., Oren, E., Polleres, A., Scicluna, J., Stollberg, M.: Web Service Modeling Ontology (WSMO), 2005.

18. Terziyan, V. Y., Kononenko, O.: Semantic web enabled web services: State-of-the-art and industrial challenges. In Proc. International Conference on Web Services (ICWS), 2003.

19. Whitworth, W. A.: Choice and Chance, with one thousand exercises. Hafner Pub. Co. New York, 1965.

Appendix A: WSDL Example Descriptions

Example Service 1: Currency Converter

```
<wsdl:definitions ... >
 <wsdl:types>
  <schema targetNamespace="http://nets.ii.uam.es/CurrencyConverter"
          xmlns="http://www.w3.org/2001/XMLSchema">
   <complexType name="CurrencyConverterRequest"> <sequence>
     <element name="amount" type="xsd:double"/>
     <element name="currencyCode1" type="xsd:string"/>
     <element name="currencyCode2" type="xsd:string"/>
    </sequence> </complexType>
   <complexType name="CurrencyConverterResponse"> <sequence>
     <element name="convertedAmount" type="xsd:double"/>
   </sequence> </complexType>
  </schema>
 </wsdl:types>
 <wsdl:message name="convertCurrencyRequest">
    <wsdl:part name="request" type="impl:CurrencyConverterRequest"/>
 </wsdl:message>
 <wsdl:message name="convertCurrencyResponse">
    <wsdl:part name="response" type="impl:CurrencyConverterResponse"/>
 </wsdl:message>
 <wsdl:portType name="CurrencyConverterPortType">
    <wsdl:operation name="convertCurrency" parameterOrder="request">
       <wsdl:input message="impl:convertCurrencyRequest"
                   name="convertCurrencyRequest"/>
       <wsdl:output message="impl:convertCurrencyResponse"
                   name="convertCurrencyResponse"/>
    </wsdl:operation>
 </wsdl:portType>
...
</wsdl:definitions>
```

Example Service 2: Trip Time Calculator

```
<wsdl:definitions ...>
 <wsdl:types>
  <schema targetNamespace="http://nets.ii.uam.es/TripTimeCalculator"
          xmlns="http://www.w3.org/2001/XMLSchema">
   <complexType name="TripTimeCalcualtorRequest"> <sequence>
     <element name="averageSpeed" type="xsd:double"/>
     <element name="departureCity" type="xsd:string"/>
     <element name="destinationCity" type="xsd:string"/>
    </sequence> </complexType>
   <complexType name="TripTimeCalculatorResponse"> <sequence>
     <element name="tripDuration" type="xsd:double"/>
   </sequence> </complexType>
  </schema>
 </wsdl:types>
 <wsdl:message name="tripTimeCalculatorRequest">
    <wsdl:part name="request" type="impl:TripTimeCalculatorRequest"/>
 </wsdl:message>
 <wsdl:message name="tripTimeCalculatorResponse">
    <wsdl:part name="response" type="impl:TripTimeCalculatorResponse"/>
 </wsdl:message>
 <wsdl:portType name="TripTimeCalculatorPortType">
    <wsdl:operation name="calculateTripTime" parameterOrder="request">
       <wsdl:input message="impl:tripTimeCalculatorRequest"
                   name="calculateTripTimeRequest"/>
       <wsdl:output message="impl:tripTimeCalculatorResponse"
                   name="calculateTripTimeResponse"/>
    </wsdl:operation>
 </wsdl:portType>
...
</wsdl:definitions>
```

Appendix B: WSMO Example Descriptions

Example Service 1: Currency Converter

```
namespace {
   _"http://nets.ii.uam.es/CurrencyConverter#",
   dc _"http://purl.org/dc/elements/1.1#",
   sample _"http://nets.ii.uam.es/SampleOntology#",
}
webService _"http://nets.ii.uam.es/CurrencyConverter"
   capability ConverterCapability
      sharedVariables {?ccode2}
      precondition
         definedBy
            ?amount memberOf sample#MoneyAmount and
               ?amount[sample#fromCurrency hasValue ?ccode1] and
            ?ccode1 memberOf sample#CurrencyCode and
            ?ccode2 memberOf sample#CurrencyCode.
      postcondition
         definedBy
            ?convertedAmount memberOf sample#MoneyAmount and
               ?convertedAmount[sample#fromCurrency hasValue ?ccode2].
```

Example Service 2: Trip Time Calculator

```
namespace {
   _"http://nets.ii.uam.es/TripTimeCalculator#",
   dc _"http://purl.org/dc/elements/1.1#",
   sample _"http://nets.ii.uam.es/SampleOntology#",
}
webService _"http://nets.ii.uam.es/TripTimeCalculator"
   capability CalculatorCapability
      precondition
         definedBy
            ?speed memberOf sample#Speed and
               ?speed[sample#inUnits hasValue "km"] and
            ?city1 memberOf sample#City and
            ?city2 memberOf sample#City.
      postcondition
         definedBy
            ?tripTime memberOf sample#Duration and
               ?tripTime[sample#inUnits hasValue "min"].
```