Ontology Mapping to support ontology-based question answering

Maria Teresa Pazienza [(*)], **Armando Stellato**[(*)],
Lina Henriksen [(+)], **Patrizia Paggio** [(+)], **Fabio Massimo Zanzotto** [(#)]

[(*)] AI Research Group, Dept. of Computer Science, Systems and Production
University of Rome, Tor Vergata
Via del Politecnico 1, 00133 Rome, Italy
{pazienza, stellato }@info.uniroma2.it

[(+)] Center for Sprogteknologi,
University of Copenhagen
Njalsgade 80
2300-DK Copenhagen S,
{lina,patrizia}@cst.dk

[(#)] DiSCo
University of Milano Bicocca
Via Bicocca degli Arcimboldi, 8
20126 Milano (Italy)
zanzotto@disco.unimib.it

Abstract

The scenario offered by the Semantic Web is a complex and impervious one: distributed service applications must rely on different information sources which reflect needs and expectations of diverse and heterogeneous cultures. In such a Babel, interaction and cooperation between these applications are still critical issues for reaching true knowledge interoperability, while integration efforts offer a viable solution to untie this Gordian knot. In this paper we report our experience in using the XeOML mapping language to harmonize two ontologies that have been developed as a support for a multilingual distributed question answering system for the Danish and Italian university domains. The different cultural aspects which affected the realization of the two information sources are discussed, along with the issues we addressed in producing a mapping between their concepts. Finally, we present overall quantitative information on the nature of the produced mapping resource.

1 Introduction

The Semantic Web (Berners-Lee et al. 2001) is offering a very interesting possibility to semantic-based NLP applications as it envisions formal semantic models available to assure interoperability between distributed software agents (Pazienza and Vindigni, 2002). A possible ontology-based Question Answering (Q&A) system in the context of the Semantic Web (SW) will use an ontology as the semantic model to interpret questions and thus retrieve the target information from relevant (possibly different) knowledge repositories.

The scenario offered by the SW is, however, a complex and heterogeneous one, where different and/or partially overlapping resources will coexist in a melting pot of distinct cultures, perspectives and representation approaches. In this situation it should be assumed that web nodes will expose information according to their own local ontology. As a consequence, Q&A systems will only be able to translate questions with respect to their underlying information sources, and those willing to interoperate in a distributed web context should tackle this diversity. A solution is provided by ontology mapping, in the form of languages, algorithms and methods that can support the translation between different knowledge resources.

In the following section, we first introduce a typical scenario where ontology mapping is necessary to achieve knowledge interoperability across distributed systems: a multi-lingual Q&A system which offers information on universities located in different countries is described in detail.

Sections 2.1 and 2.2 offer a deeper look on the ontology mapping task, providing details on the two ontologies, the cultural differences which characterize the two domains they represent and the requirements which a mapping language must match to support the task they have been developed for.

In Section 3, the adoption of the language XeOML for mapping the two ontologies has been discussed, followed by details of our case study: how the mapping task has been performed, which methodology has been applied, a description of specific examples which show the use of different mapping primitives and some statistics which provide qualitative and quantitative information on the coverage of the two domains.

2 The MOSES scenario

Our approach to ontology mapping is being developed in the context of the EU project MOSES[1] IST-2001-37244. The project's overall objective is to develop an ontology-based

[1] MOSES is a cooperative project under the 5th Framework Programme. The project partners are FINSA Consulting, MONDECA, Centre for Language Technology and the Faculty of Humanities from the University of Copenhagen, University of Roma Tre, University of Roma Tor Vergata and ParaBotS.

methodology to create, maintain, search and adapt semantically structured Web contents in a federation of sites (Atzeni et al, 2004) (Paggio et al. 2004). As a test bed the project is developing an agent-based knowledge management system and an ontology-based search engine that will both accept questions and produce answers in natural language for the Web sites of the two European universities of Roma Tre and Copenhagen. The specific purpose of ontology mapping is to support question answering across the two nodes, each of which is equipped with its own conceptualization of the domain.

In this scenario, a user submits questions to the system from one of the two sites in the language pertaining to this site, and indicates while doing so whether the question is to be interpreted as a single-node one to be processed locally or a federated one. In the case of a federated scenario, access to all the nodes in the federation must be supported. After having analysed the input in terms of the local analyser, the local content agent will issue a mapping request to the agents responsible for the remote nodes to get the query mapped onto the remote ontologies and processed by the remote content agents. The content retrieved from the remote nodes will then be combined with the relevant local content and presented to the user in a convenient and understandable way, for instance as shown – in English for readability's sake – in Figure 1 below.

Q: Who teaches history?

The following answers are found on the KU site:

"Lars Hansen underviser i historie"

The following answers are found on the ROMA III site:

"Paolo Rossi insegna storia dell'arte"

Figure 1: An interaction in the federated scenario

The federated scenario is illustrated in Figure 2 by means of three sites: the Danish site S_D, the Italian site S_I, and an additional site S_X. In this example, the user communicates with the system via the Danish node, in other words the question is in Danish, and the answer is constructed and presented to the user by the Danish answer generator, although the data found can in principle originate from all three sites. The figure also illustrates the fact that each node in the federation is equipped with its own ontology and knowledge base. Therefore, for a query to be propagated across several nodes, a mapping between the ontologies involved is necessary.

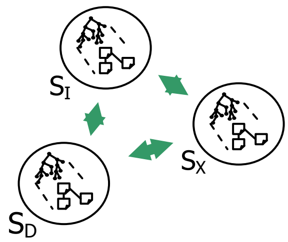

Figure 2: Querying in a federated scenario

2.1 The MOSES ontologies

The ontologies underlying the Danish and Italian nodes of the MOSES test bed have been built with the "university ontology" from the DAML-OIL ontology library as a starting point. The classes and associations of this ontology cover in fact, at least at a high level, most of the relevant concepts relevant to the MOSES scenario (i.e. *University Staff and Students, Courses* and *Research).* Several changes and extensions, however, were necessary to cover the two national university sub-systems (i.e. Italian and Danish). Coverage figures of the resulting ontologies are shown in Table 2. Instances have been created by the project's user groups by downloading them from the respective sites' databases as well as by manually extracting data from the Web pages.

The first challenge deriving from having these two separate ontologies for the same domain is the language. Classes and relations in the two ontologies are in fact labelled in different languages. Another challenge comes from the structural differences: not all the nodes in one ontology are represented also in the other and vice-versa. Finally, domain relations are treated differently in the two ontologies. In the Italian one, all relations are binary in keeping with the original DAML-OIL model, whereas the Danish ontology makes use of n-nary relations in the spirit of the Topic Maps formalism (Garshol, 2003) used by the MOSES content management system. Note that the term "association" is used in Topic Maps instead of relation, so from here on we will use this term to refer to the relations in the two ontologies.

2.2 Two domains or one extended domain?

One issue that is primarily addressed when building a harmonization of two knowledge resources is related to the analysis of the overlap that exists between the domains they describe. Differences should be discussed and analyzed at a general level to arrive at overall considerations on the nature of two ontologies and of their modeled domains, but also at single concept level.

Many approaches to ontology mapping rate the similarity of two concepts upon the overlap that exists between their projection over the same domain. Concept extension is thus considered as an ideal metric for similarity measures: in (Doan et al., 2002) a probabilistic approach adopts the Jaccard measure (Padro, 1998) to compute the joint distribution of two concepts, i.e. the probability that a given element of the domain is included in the extension of both the concepts. C-OWL (Bouquet et al., 2003) follows an "exact" approach providing semantics for equivalence, inclusion, orthogonality, and overlap, again described in terms of coverage of domain objects. Other approaches, like MAFRA (Maedche, 2002) do not address the problem of describing the differences which bind concepts to their domain interpretation, trying instead to define mappings in terms of pre and post conditions which ensure conceptual equivalence. We believe, nevertheless, that some discrepancies in the represented domains, as well as the peculiarities of the task to be supported, may require the assessment of relations which go beyond the idea of "domain overlap".

In the MOSES scenario, it is a matter of interpretation whether the two university ontologies describe two distinct domains or constitute different conceptualizations of the same. If we choose to see them as modeling one extended domain, characterizing as a minimum the two universities taken together, or perhaps all the Italian and the Danish universities, and maximally the set of all the universities, we have a model on which to define mapping equivalence. In the extended domain, all instances of the ***University*** class in the Italian ontology are valid instances of the **Universitet** class in the Danish, and vice-versa: the two classes can be considered equivalent. However, equivalence does not suffice if we want the mapping to support question answering across the two ontologies. For instance, if we consider questions involving the Italian class **Professore Associato** and look for an equivalent concept in the Danish academic career, we will not be able to provide any relevant data. However, it seems reasonable to expect that **Professore Associato** is related to the Danish **Lektor**, which describes a similar type of position in the Danish system. To support question answering across ontologies, we need a mapping language that is flexible enough to allow for the definition of equivalence as well as similarity relations, and their use in simple as well as complex mappings between different ontology elements.

3 Mapping Ontologies with XeOML: A language for Ontology Mapping

XeOML (Pazienza et al., 2004) is a language for describing ontology mappings which has been developed at the University of Rome, Tor Vergata.

The principle beyond XeOML is to define a layered view in which the complex logical-algebraic aspects of the ontology mapping task coexist (but are neatly separated) with the semantics which describe the nature and the quality of the identified correspondences.

In XeOML a neutral ontology structure is being adopted to achieve interoperability between different representation languages which share an object oriented approach to knowledge modeling. Abstract definitions for *instances*, *classes*, *properties* and *associations*, together with the description of the mapping primitives that operate on them, are thus provided in a core language, which can then be extended to fit on specific representation primitives and to support the different semantic aspects of the mapping task it is being applied to.

The idea underlying the extensible definitions of element types to different representation formalisms, is that a mediation activity involving two agents, requires them to be only proficient about the knowledge model adopted to express their underlying ontological resource while not necessarily being able to understand the model owned by the interlocutor.

XeOML is thus not intended as a set of mapping relations with well defined semantics, as for C-OWL, or of primitives for operating data transformation, like the ones adopted in the MAFRA System. XeOML is an open environment providing:

- the structure of a mapping document
- a set of constructs for specifying complex concepts (see below for details) to be defined inside the scope of specific mappings
- an extensible mechanism for including existent semantic relations (like those from OWL/C-OWL) or newly defined ones.

3.1 XeOML Language characteristics

The XeOML[2] core language described above is specified by an XML schema, *AbstractMapping*, which defines the structure of a mapping document (MD) and declares the set of elements that populate an ontology.

To obtain a concrete schema for instantiating a MD between ontologies, two extensions to the AbstractMapping schema must be provided:

[2] XeOML related material is freely available for download at: http://ai-nlp.info.uniroma2.it/xeoml

- an *Ontology Elements Definition* schema extension, which accounts for specification of ontological elements according to a given representation language; possibly, two extensions may be required if the ontologies to be mapped are represented through different languages
- a *Mappings Definition* schema extension, where qualitative (and possibly quantitative) measures for rating the nature of the asserted mapping may be provided.

More semantically declarative information may thus be plugged to the main schema in the form of XML Schema extensions, which reflect different perspectives and approaches to the mapping process and/or heterogeneous knowledge representation styles; the core schema, together with its extensions, forms a concrete MD schema.

Mappings are conceptually divided into two categories: *Simple Mappings*, which define one-to-one relations between ontology elements of the same type, and *Complex Mappings,* which link structures of heterogeneous objects from the two ontologies. In a XeOML document, a <mapping> entity may represent either a simple mapping on its own or contain a single concept and a reference to a complex mapping. This way, the <mapping> entity acts as an index for the mapping function.

Simple mappings between classes, instances or association roles include *Equivalence*, *Similarity* and *TypeMismatch* (SuperClass, SubClass, etc.).

Complex mappings vary in the nature of the concepts involved and in the operations that are applied over them, including:

- *restrictions* on classes/associations operated on the range of their attributes
- *aggregations* (on an extensional basis) of multiple classes/associations
- *transformations* between heterogeneous structures of objects (properties vs associations+roles, classes vs instances)
- *join* of associations upon common roles

3.2 XeOML extensions implementation inside MOSES: the “extended domain” interpretation

Since the two ontologies adopt the same knowledge model, a single implementation of the *Ontology Elements* schema has been produced, containing the appropriate definitions for Topic Maps concepts.

The *Mappings Definition* schema is instead the *locus* for specifying the semantics which are applied to the mapped complex structures defined in the main schema. As the two ontologies are describing virtually the same domain – apart from differences which reflect the cultural aspects of the involved countries – we modeled the mapping relations considering an extended perspective of the domain, which includes objects represented in both the ontologies, and which we call *the extended domain*.

The mapping relations between classes have thus been defined to express mutual coverage of the newly consolidated domain, in terms of extensional equivalence (*DenotationalEquivalence*) and inclusion (*SubClassOf*, *SuperClassOf)*, or to address disjoint set of objects which are supposed to be bound by similar roles in their respective (original) domain (*Similarity*).

Under this perspective, the previously discussed Italian class **ProfessoreAssociato** and the Danish **Lektor** are considered *Similar*, because they share no overlap over the extended domain, though addressing similar roles inside the Danish and Italian academic careers. The Italian class **Professore** and the Danish **Professorat**, are instead considered *Denotationally Equivalent* because, in the extended domain, they both comprehend the union of their extensions from their original domains: a Danish **Lektor** is axiomatically a subclass of **Professorat** in the Danish ontology**,** and is also considered a **Professore** according to the interpretation of this class over the extended domain.

The same principle has been applied when considering instances: in this case the choice is between a strict equivalence (i.e. the interpretation of two instances is referring the same domain object) or again, an affinity of roles inside their own domains.

Mappings involving Relational Properties and Associations have been described in terms of differences in their ranges/roles.

No data transformation for properties has been considered, as the wide majority of the properties in the two ontologies are relational properties, while the relatively small set of data properties (e.g. e-mail) uses analogous formats for representing their content.

4 Developing the case study: a concrete mapping

In MOSES a concrete XeOML mapping instance has been created manually between the Danish and the Italian ontologies with at least two goals. First of all, the mapping is used in the agent platform to support communication between content agents of different nodes. But another important goal is to support the development and testing of algorithms for automatic mapping by establishing a gold standard against which such algorithms can be evaluated. Automatic mapping is not the focus in this paper, and will not be discussed any further. As for the exploitation of the mapping in agent communication, we provide below a discussion of the mapping relations involved in a specific example, but leave out, still in keeping with the focus of the paper, an account of the model used in the agent platform.

In order to achieve a somewhat balanced result, the two groups participating in the mapping task – RTV and CST – defined each a partial mapping to start with, and each group considered its own ontology as the left-hand side of the mapping relations to ensure maximum coverage. The two independent results were then compared, divergences were discussed and solved, and a final mapping merging the two independent ones was produced. In the final result, which can be ispected from the XeOML site at http://ai-nlp.info.uniroma2.it/xeoml/, only a subset of the mapping relations available in XeOML have been adopted. For instance, only one complex mapping relation – *ClassAggregation_ComplexMap* – is used for classes. However, a fair number of mapping relations have been explored.

4.1 Applying the mapping: an example

Let us now assume that the user, while accessing the Danish site, asks the question *Hvem underviser i historie?* (Who teaches history?), and indicates by clicking on the relevant menu button that it is a federated question. The input will be analysed by the Danish linguistic analyser, which will produce a query asking for all the 'course' (**Kursus**) topics participating in two associations: *i.* a 'studySubject' (**StudieværkEmne**) association where the 'subject' (**Emne**) role is specified to be the instance 'history' (***historie***); and *ii.* a 'courseOffer' (**KursusUdbud**) association, which in addition to the 'course' role also includes roles played by the classes 'teacher' (**Lærerstab**) and 'universityOrgan' (**Universitetsorgan**). In order to retrieve relevant information from the Italian node, local concepts and associations must be mapped into the Italian conceptual language, and a new query reconstructed around these. Table 1 shows all the relevant mapping relations for this example.

Two different mappings of the Danish **KursusUdbud** into the Italian ontology are possible in that either **teacherOf** or **teachingAssistantOf** can be aggregated and bounded to **offersCourse** to create a valid mapping for the three roles involved in **KursusUdbud**, and therefore of the association as a whole. Finally, a mapping is also available from the **StudieværkEmne** association to the Italian **workSubject**. In conclusion, all the arguments playing a role in the query can be mapped onto meaningful concepts in the Italian ontology, although the relations are not all equally reliable (with equivalence being of course the strongest). In other cases, only part of the information represented in the input query may be expressible in terms of the target ontology.

One final observation concerns the mapping at instance level, in our example from the Danish ***filmhistorie*** to the Italian ***storia del film***. Instances have been left out of the mapping established between the two ontologies. This seems a reasonable choice, first of all because a mapping in many cases (named individuals) will not make sense, and secondly because, in cases such as names of disciplines where it does make sense, the task is an open-ended one which falls beyond the scope of the project. Lists of some relevant subjects and their translations have been created as an initial practical solution.

Danish Associations/ Classes	**Italian Associations/ Classes**	**Mapping relations**
KursusUdbud	**teacherOf, offersCourse**	**Aggregation ParticipantsIS-AMismatch**
Universitetsorgan	University	SuperClass (Universitetsorgan, Universitet) Equivalence (Universitet, University)
Kursus	Course	Similar (Kursus, Course)
Lærerstab	Faculty	Super (Lærerstab, Faculty)
KursusUdbud	**teachingAssistantOf, offersCourse**	**Aggregation ParticipantsIS-AMismatch**
Universitetsorgan	University	SuperClass (Universitetsorgan, Universitet) Equivalence (Universitet, University)
Kursus	Course	Similarity (Kursus, Course)
Lærerstab	teachingAssistant	SuperClass (Lærerstab, Undervisningsassistent) Similarity (Undervisningsassistent, TeachingAssistent)
StudieværkEmne	**workSubject**	**ParticipantsIS-Amismatch**
Studieværk	Work	SubClass (Studieværk, Værk), Similarity(Værk, Work)
Emne	Subject	Similarity (Emne, Subject)

Table 1: Mapping relations for federated question

4.2 Completeness of the mapping document

Our approach at ontology mapping aims at reducing the cost related to manual production of the mapping relations. Inference rules for deriving relations which are not explicitly mentioned in the MD have thus been defined. These rules include a back-off strategy for mapping properties (which first looks for a contextualized mapping of a property, i.e. when the property is considered attached to a particular class, and then goes back to the general mapping for that property) and rules for recursively exploring the ontology taxonomies to derive new mappings for classes.

We consider a notion of completeness based on a committed interpretation of the two mapped ontologies. A concept from an ontology *needs to be mapped* if any of the semantic relations defined in the applied XeOML concrete schema may relate it with a concept from the other ontology. So, a MD is complete iff all the concepts of the two ontologies which *need to be mapped* have been explicitly mapped through the MD or through a relation that can be inferred from the existing ones.

4.3 Statistics of the mapping document

We report in table 2 some statistical information on the released mapping document. We did not produce a gold standard of the mapping document for computing completeness of the mappings between the two resources. There are different reasons for this choice: at the moment no standard for mapping evaluation is being widely adopted, and the reason is probably that a simple precision/recall measure, though reporting the number of correctly mapped concepts versus the wrongly assigned and the missed ones, would not be quite informative of the quality of the produced (explicitly or by inference) mappings; at the same time, more complex measures should depend on the nature of the different involved relations and should thus vary from experiment to experiment.

In our experiment, two MDs have been produced after the work of two independently employed human annotators, which were previously trained on the task; on a second step, their MDs have been compared and their differences judged by a third annotator, which then released the final MD. In this sense, the MD is correct "by definition", *recall* is automatically 100% (as a consequence of our methodology, where a concept needs to be mapped explicitly only if its related mapping is not entailed by one of the already asserted ones) and *precision* is guaranteed by correct entailments between the explicit mappings.

Therefore, rather than expressing percentages, we neutrally report the total number of concepts in the two ontologies (the **Source** entries), the number of explicitly mapped concepts (**MD**), and the number of concepts mappable through implied relationships, to get an idea of the progressive trend in the coverage offered by the XeOML inferential engine.

In particular, the **D-Imp** (*Descending Implicated*) entry refers to those concepts that may have no direct counterpart in the other ontology, but that at least subsume one or more mappable concepts. Concepts mapped through an *Ascending Implicated* (those under the **A-Imp**) are instead subsumed by some concept which can be mapped. Finally, the **Imp** column refers to those concepts for which there exists at least one of the previously described mappings (**A-Imp** and **D-Imp** are not disjoint).

Descending Implicated bring about a valid and consistent answer for the Q&A system, as the retrieved concepts extensionally identify correct instances of the desired one, while *Ascending Implicated* will generate an imprecise answer. For this reason, our mapping algorithm first looks for a direct match, then carries out a downward search and, in case of a fail, reports *Ascending Implicated.* In many ontologies, this may lead to the proliferation of dangerous generalizations. To limit this problem, a lower bound for the precision of an *Ascending Implicated* can be set by defining a set of "not allowed generalizers", i.e. classes which are forbidden in an upward search. In our experiment, we had three root classes for both the ontologies, which have been inherited from the general Topic Map ontology: **Subject**, **Document** and **Association**. We included **Document**, **Association** plus **Universal** and **Individual** (the two sibling subclasses of **Subject**) into the set of not allowed generalizers, and calculated the results on

	Italian Ontology						**Danish Ontology**				
	Source	**MD**	**D-Imp**	**A-Imp**	**Imp**		**Source**	**MD**	**D-Imp**	**A-Imp**	**Imp**
Classes (*Subject*)	94	50	55	86	89		152	54	67	146	147
Classes (*Document*)	48	17	19	27	29		22	17	17	20	20
Associations	47	24	X	X	X		36	18	X	X	X
ObjectProperties	95	52	X	X	X		58	37	X	X	X
DatatypeProperties	18	4	X	X	X		10	3	X	X	X

Table 2: Statistics of the MOSES mapping experiment

this basis. This choice led to a good compromise between domain coverage and quality of the implied mappings: for example, the 5 classes which have been left out from the set of Italian Subjects, include concepts like “Artificial Agent”, “Function” and “Organism” which in fact represent abstract ideas, vague concepts or ad hoc constructs that can hardly be related to anything in the Danish ontology. We left an ‘X’ under the boxes for implied property mappings, as in our case study there was no property hierarchy. The same holds for Associations, as their corresponding classes are reported in a flat set under the **Association** class.

5 Conclusion and future research

In a strongly distributed context like the one predicted by the Semantic Web Vision, it appears unrealistic to assume that every web node be based on a centralized knowledge repository.

What will really happen is that many different heterogeneous ontologies with overlapping domains will be developed and exploited by several partners of the communication.

The standardization of languages for representing knowledge is, in this sense, only a first achieved step, which must necessarily be followed by a multitude of proposals and efforts aiming at reaching semantic consensus among distributed knowledge resources.

In this paper, we have presented our experience in mapping two ontologies that, though describing overlapping domains, showed heavy structural and semantic differences which needed to be harmonized. The XeOML language and inference rules have proved to be a valid mean for supporting ontology mapping in situations which may vary depending on the task and domain the ontologies are dedicated to.

Future research work will go in the direction of further reducing the manual effort which must be dedicated to the mapping task, developing tools for semi-automatic production of mapping documents, featuring improved inferential capabilities and integrated visualization of the ontologies and their associated mappings. Also, though maintaining our neutral and extensible approach, we are considering moving the abstract XeOML schema from pure XML to OWL, thus easing the integration of existent OWL mapping properties (or those from its extensions, like for C-OWL), which would bring stronger semantics to our generic mapping primitives.

References

P. Atzeni, R. Basili, D. H. Hansen, P. Missier, P. Paggio, M. T. Pazienza, F. M. Zanzotto. 2004. *Ontology-based question answering in a federation of university sites: the MOSES case study.* 9th International Conference on Applications of Natural Language to Information Systems (NLDB'04) Manchester (United Kingdom).

T. Berners-Lee, J. Hendler, and O. Lassila. 2001. *The Semantic Web: A new form of Web con-tent that is meaningful to computers will unleash a revolution of new possibilities*. Scientific American.

P. Bouquet, F. Giunchiglia, F. van Harmelen, L. Serafini and H. Stuckenschmidt. 2003. *C-OWL: Contextualizing ontologies*. In K. Sekara and J. Mylopoulis, editors. Proceedings of the Second International Semantic Web Conference, number 2870 in Lecture Notes in Computer Science, pages 164–179. Springer Verlag.

A. Doan, J. Madhavan, P. Domingos and A. Halevy. 2002. *Learning to map between ontologies on the semantic web*. In The Eleventh International WWW Conference, Hawaii, US.

A. Maedche, B. Motik, N. Silva and R Volz. 2002. *Mafra a mapping framework for distributed ontologies*. In Proceedings of the 13th European Conference on Knowledge Engineering and Knowledge Management EKAW-2002, Madrid Spain

L. Padro. 1998. *A Hybrid Environment for Syntax-Semantic Tagging*.

P. Paggio, D. H. Hansen, R. Basili, M. T. Pazienza, F. M. Zanzotto. 2004. *Ontology-based question analysis in a multilingual environment*: *the MOSES case study.* Workshop OntoLex 2004, held jointly with LREC 2004 LISBON, Portugal.

M. T. Pazienza, A. Stellato, M. Vindigni, F. M. Zanzotto. 2004. *XeOML: An XML-based extensible Ontology Mapping Language.* Workshop on Meaning Coordination and Negotiation, held in conjunction with 3rd International Semantic Web Conference (ISWC-2004) Hiroshima, Japan.

M.T. Pazienza, M. Vindigni. 2002 *Agent based Ontological Mediation in IE Systems*. Information Extraction in the Web Era: Natural Language Communication for Knowledge Acquisition and Intelligent Information Agents, papers from SCIE 2002, Frascati (Rome, Italy), Springer 2003.