

Towards Agent-based Cross-lingual Interoperability of Distributed Lexical Resources

Claudia Soria^{*}

Maurizio Tesconi[°]

Andrea Marchetti[°]

Francesca Bertagna^{*} Monica Monachini^{*}

Chu-Ren Huang[§] Nicoletta Calzolari^{*}

^{*}CNR-ILC and [°]CNR-IIT
Via Moruzzi 1, 56024 Pisa
Italy

[§]Academia Sinica
Nankang, Taipei
Taiwan

{firstname.lastname@ilc.cnr.it}
{firstname.lastname@iit.cnr.it}

churen@gate.sinica.edu.tw

Abstract

In this paper we present an application fostering the integration and interoperability of computational lexicons, focusing on the particular case of mutual linking and cross-lingual enrichment of two wordnets, the ItalWordNet and Sinica BOW lexicons. This is intended as a case-study investigating the needs and requirements of semi-automatic integration and interoperability of lexical resources.

1 Introduction

In this paper we present an application fostering the integration and interoperability of computational lexicons, focusing on the particular case of mutual linking and cross-lingual enrichment of two wordnets. The development of this application is intended as a case-study and a test-bed for trying out needs and requirements posed by the challenge of semi-automatic integration and enrichment of practical, large-scale multilingual lexicons for use in computer applications. While a number of lexicons already exist, few of them are practically useful, either since they are not sufficiently broad or because they don't cover the necessary level of detailed information. Moreover, multilingual language resources are not as widely available and are very costly to construct: the work process for manual development of new lexical resources or for tailoring existing ones is too expensive in terms of effort and time to be practically attractive.

The need of ever growing lexical resources for effective multilingual content processing has urged the language resource community to call for a radical change in the perspective of language resource creation and maintenance and the design of a “new generation” of LRs: from static, closed and locally developed resources to shared and distributed language *services*, based on open content interoperability standards. This has often been called a “change in paradigm” (in the sense of Kuhn, see Calzolari and Soria, 2005; Calzolari 2006). Leaving aside the tantalizing task of building on-site resources, the new paradigm depicts a scenario where lexical resources are cooperatively built as the result of controlled cooperation of different agents, adopting the paradigm of accumulation of knowledge so successful in more mature disciplines, such as biology and physics (Calzolari, 2006).

According to this view (or, better, this *vision*), different lexical resources reside over distributed places and can not only be accessed but *choreographed* by agents presiding the actions that can be executed over them. This implies the ability to build on each other achievements, to merge results, and to have them accessible to various systems and applications.

At the same time, there is another argument in favor of distributed lexical resources: language resources, lexicons included, are inherently distributed because of the diversity of languages distributed over the world. It is not only natural that language resources to be developed and maintained in their native environment. Since language evolves and changes over time, it is not possible to describe the current state of the lan-

guage away from where the language is spoken. Lastly, the vast range of diversity of languages also makes it impossible to have one single universal centralized resource, or even a centralized repository of resources.

Although the paradigm of distributed and interoperable lexical resources has largely been discussed and invoked, very little has been made in comparison for the development of new methods and techniques for its practical realization. Some initial steps are made to design frameworks enabling inter-lexica access, search, integration and operability. An example is the Lexus tool (Kemps-Snijders et al., 2006), based on the Lexical Markup Framework (Romary et al., 2006), that goes in the direction of managing the exchange of data among large-scale lexical resources. A similar tool, but more tailored to the collaborative creation of lexicons for endangered language, is SHAWEL (Gulrajani and Harrison, 2002). However, the general impression is that little has been made towards the development of new methods and techniques for attaining a concrete interoperability among lexical resources. Admittedly, this is a long-term scenario requiring the contribution of many different actors and initiatives (among which we only mention standardisation, distribution and international cooperation).

Nevertheless, the intent of our project is to contribute to fill in this gap, by exploring in a controlled way the requirement and implications posed by new generation multilingual lexical resources. The paper is organized as follows: section 2 describes the general architectural design of our project; section 3 describes the module taking care of cross-lingual integration of lexical resources, by also presenting a case-study involving an Italian and Chinese lexicons. Finally, section 4 presents our considerations and lessons learned on the basis of this exploratory testing.

2 An Architecture for Integrating Lexical Resources

LeXFlow (Soria et al., 2006) was developed having in mind the long-term goal of lexical resource interoperability. In a sense, LeXFlow is intended as a proof of concept attempting to make the *vision* of an infrastructure for access and sharing of linguistic resources more tangible.

LeXFlow is an adaptation to computational lexicons of XFlow, a cooperative web application for the management of document workflows

(DW, Marchetti et al., 2005). A DW can be seen as a process of cooperative authoring where a document can be the goal of the process or just a side effect of the cooperation. Through a DW, a document life-cycle is tracked and supervised, continually providing control over the actions leading to document compilation. In this environment a document travels among agents who essentially carry out the pipeline receive-process-send activity.

There are two types of agents: *external agents* are human or software actors performing activities dependent from the particular Document Workflow Type; *internal agents* are software actors providing general-purpose activities useful for many DWTs and, for this reason, implemented directly into the system. Internal agents perform general functionalities such as creating/converting a document belonging to a particular DW, populating it with some initial data, duplicating a document to be sent to multiple agents, splitting a document and sending portions of information to different agents, merging duplicated documents coming from multiple agents, aggregating fragments, and finally terminating operations over the document. External agents basically *execute* some processing using the document content and possibly other data; for instance, accessing an external database or launching an application.

LeXFlow was born by tailoring XFlow to management of lexical entries; in doing so, we have assumed that each lexical entry can be modelled as a document instance, whose behaviour can be formally specified by means of a *lexical workflow type* (LWT). A LWT describes the life-cycle of a lexical entry, the agents allowed to act over it, the actions to be performed by the agents, and the order in which the actions are to be executed. Embracing the view of cooperative workflows, agents can have different rights or views over the same entry: this nicely suits the needs of lexicographic work, where we can define different roles (such as encoder, annotator, validator) that can be played by either human or software agents. Other software modules can be inserted in the flow, such as an automatic acquirer of information from corpora or from the web. Moreover, deriving from a tool designed for the cooperation of agents, LeXFlow allows to manage workflows where the different agents can reside over distributed places.

LeXFlow thus inherits from XFlow the general design and architecture, and can be considered as a specialized version of it through design

of specific Lexical Workflow Types and plug-in of dedicated external software agents. In the next section we briefly illustrate a particular Lexical Workflow Type and the external software agents developed for the purpose of integrating different lexicons belonging to the same language. Since it allows the independent and coordinated sharing of actions over portions of lexicons, LeXFlow naturally lends itself as a tool for the management of distributed lexical resources.

Due to its versatility, LeXFlow is both a general framework where ideas on automatic lexical resource integration can be tested and an infrastructure for proving new methods for cooperation among lexicon experts.

2.1 Using LeXFlow for Lexicon Enrichment

In previous work (Soria et al., 2006), the LeXFlow framework has been tested for integration of lexicons with differently conceived lexical architectures and diverging formats. It was shown how interoperability is possible between two Italian lexicons from the SIMPLE and WordNet families, respectively, namely the SIMPLE/CLIPS (Ruimy et al., 2003) and Ital-WordNet (Roventini et al., 2003) lexicons.

In particular, a Lexical Workflow Type was designed where the two different monolingual semantic lexicons interact by reciprocally enriching themselves and moreover integrate information coming from corpora. This LWT, called “lexicon augmentation”, explicitly addresses dynamic augmentation of semantic lexicons. In this scenario, an entry of a lexicon A becomes enriched via basically two steps. First, by virtue of being mapped onto a corresponding entry belonging to a lexicon B , the entry $_A$ inherits the semantic relations available in the mapped entry $_B$. Second, by resorting to an automatic application that acquires information about semantic relations from corpora, the acquired relations are integrated into the entry and proposed to the human encoder.

An overall picture of the flow is shown in Figure 1, illustrating the different agents participating in the flow. Rectangles represent human actors over the entries, while the other figures symbolize software agents: ovals are internal agents and octagons external ones. The two external agents involved in this flow are the “relation calculator” and the “corpora extractor”. The first is responsible for the mapping between the sets of semantic relations used by the different lexicons. The “corpora extractor” module invokes an application that acquires information

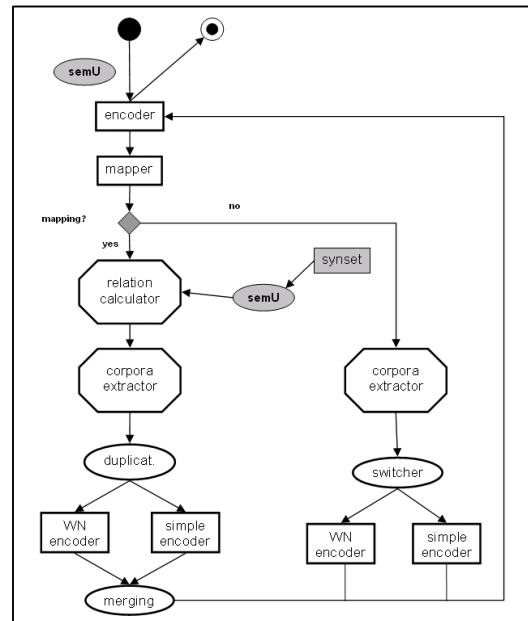


Figure 1. Lexicons Augmentation Workflow Type.

about *part-of* relations by identifying syntactic constructions in a vast Italian corpus. It then takes care of creating the appropriate candidate semantic relations for each lemma that is proposed by the application.

A prototype of LeXFlow has been implemented with an extensive use of XML technologies (XML Schema, XSLT, XPath, XForms, SVG) and open-source tools (Cocoon, Tomcat, MySQL). It is a web-based application where human agents interact with the system through an XForms browser that displays the document to process as a web form whereas software agents interact with the system via web services.

3 Multilingual WN Service

In the Section above we have illustrated the general architecture of LeXFlow and showed how a Lexical Workflow Type can be implemented in order to enrich already existing lexicons belonging to the same language but realizing different models of lexicon encoding. In this section we move to a cross-lingual perspective of lexicon integration. We present a module that similarly addresses the issue of lexicon augmentation or enrichment focusing on mutual enrichment of two wordnets in *different languages and residing at different sites*.

This module, named “multilingual WN Service” is responsible for the *automatic cross-lingual fertilization* of lexicons having a Word-

Net-like structure. Put it very simply, the idea behind this module is that a monolingual wordnet can be enriched by accessing the semantic information encoded in corresponding entries of other monolingual wordnets.

Since each entry in the monolingual lexicons is linked to the Interlingual Index (ILI, cf. Section 3.1), a synset of a WN(A) is indirectly linked to another synset in another WN(B). On the basis of this correspondence, a synset(A) can be enriched by importing the relations that the corresponding synset(B) holds with other synsets(B), and vice-versa. Moreover, the enrichment of WN(A) will not only import the relations found in WN(B), but it will also propose target synsets in the language(A) on the basis of those found in language(B).

The various WN lexicons reside over distributed servers and can be queried through web service interfaces. The overall architecture for multilingual wordnet service is depicted in Figure 2.

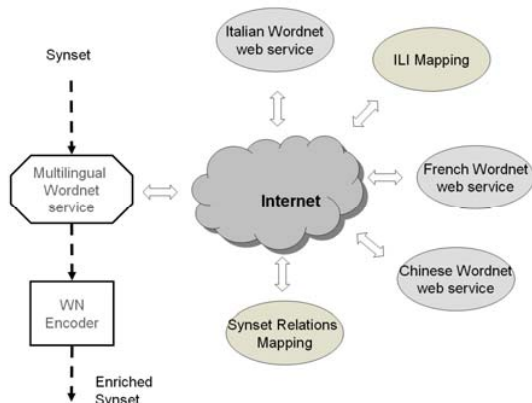


Figure 2. Multilingual Wordnet Service Architecture.

Put in the framework of the general LeXFlow architecture, the Multilingual wordnet Service can be seen as an additional external software agent that can be added to the augmentation workflow or included in other types of lexical flows. For instance, it can be used not only to enrich a monolingual lexicon but to bootstrap a bilingual lexicon.

3.1 Linking Lexicons through the ILI

The entire mechanism of the Multilingual WN Service is based on the exploitation of Interlingual Index (Peters et al., 1998), an unstructured version of WordNet used in EuroWordNet (Vossen et al., 1998) to link wordnets of different languages; each synset in the language-specific wordnet is linked to at least one record of the ILI

by means of a set of equivalence relations (among which the most important is the EQ_SYNONYM, that expresses a total, perfect equivalence between two synsets).

Figure 6 describes the schema of a WN lexical entry. Under the root “synset” we find both internal relations (“synset relations”) and ILI Relations, which link to ILI synsets.

Figure 3 shows the role played by the ILI as set of pivot nodes allowing the linkage between concepts belonging to different wordnets.

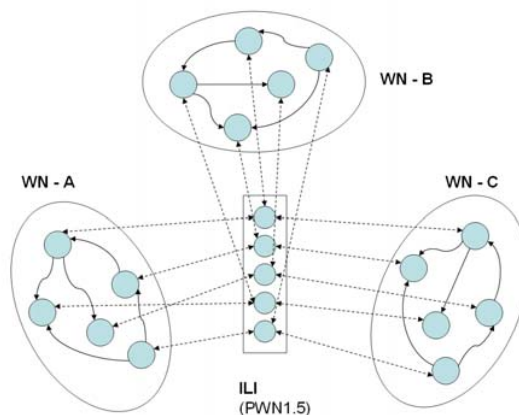


Figure 3. Interlingual Linking of Language-specific Synsets.

In the Multilingual WN Service, only equivalence relations of type EQ_SYNONYM and EQ_NEAR_SYNONYM have been taken into account, being them the ones used to represent a translation of concepts and also because they are the most exploited (for example, in IWN, they cover about the 60% of the encoded equivalence relations). The EQ_SYNONYM relation is used to realize the one-to-one mapping between the language-specific synset and the ILI, while multiple EQ_NEAR_SYNONYM relations (because of their nature) might be encoded to link a single language-specific synset to more than one ILI record. In Figure 4 we represented the possible relevant combinations of equivalence relations that can realize the mapping between synsets belonging to two languages. In all the four cases, a synset “a” is linked via the ILI record to a synset “b” but a specific procedure has been foreseen in order to calculate different “plausibility scores” to each situation. The procedure relies on different rates assigned to the two equivalence relations (rate “1” to EQ_NEAR_SYNONYM relation and rate “0” to the EQ_SYNONYM). In this way we can distinguish the four cases by assigning respectively a weight of “0”, “1”, “1” and “2”.




Figure 4. Possible Combinations of Relations between two Lexicons A and B and the ILI.

The ILI is a quite powerful yet simple method to link concepts across the many lexicons belonging to the *WordNet-family*. Unfortunately, no version of the ILI can be considered a standard and often the various lexicons exploit different version of WordNet as ILI¹. This is a problem that is handled at web-service level, by incorporating the conversion tables provided by (Daudé et al., 2001). In this way, the use of different versions of WN does not have to be taken into consideration by the user who accesses the system but it is something that is resolved by the system itself². This is why the version of the ILI is a parameter of the query to web service (see Section below).

3.2 Description of the Procedure

On the basis of ILI linking, a synset can be enriched by importing the relations contained in the corresponding synsets belonging to another wordnet.

In the procedure adopted, the enrichment is performed on a synset-by-synset basis. In other words, a certain synset is selected from a wordnet resource, say WN(A). The cross-lingual module identifies the corresponding ILI synset, on the basis of the information encoded in the synset. It then sends a query to the WN(B) web service providing the ID of ILI synset together with the ILI version of the starting WN. The WN(B) web service returns the synset(s) corresponding to the WN(A) synset, together with reliability scores. If WN(B) is based on a different ILI version, it can carry out the mapping between ILI versions (for instance by querying the ILI mapping web service). The cross-lingual module then analyzes the synset relations encoded in the

WN(B) synset and for each of them creates a new synset relation for the WN(A) synset.

If the queried wordnets do not use the same set of synset relations, the module must take care of the mapping between different relation sets. In our case-study no mapping was needed, since the two sets were completely equivalent.

Each new relation is obtained by substituting the target WN(B) synset with the corresponding synset WN(A), which again is found by querying back the WN(A) web service (all these steps through the ILI). The procedure is formally defined by the following formula:

```

Let  $a_j \in A$ 
Let  $Ba_j = \{b_i \mid b_i \in B \text{ and } (b_i \text{ ILI } a_j)\}$ 
 $\forall b_i \in Ba_j$ 
  Let  $R_i = \{b_i r_k b_p \mid b_i, b_p \in B \text{ and } (r_k \in R_A \cap R_B)\}$ 
   $\forall b_i r_k b_p \in R_i$ 
    Let  $Ab_p = \{a_i \mid a_i \in A \text{ and } (a_i \text{ ILI } b_p)\}$ 
     $\forall a_t \in Ab_p$ 
       $a_j r_k a_t$  is a candidate relation

Legenda:
A, B lexicons
 $a_j, b_i$  synsets
 $a_j r_p a_i$  synset relation  $r_p$  between  $a_j$  and  $a_i$ 
 $b_i \text{ ILI } a_j$   $b_i$  is connected by ILI with  $a_j$ 
 $R_A, R_B$  relation space of lexicons B
 $R_A \cap R_B$  the common relation space of B and A
  
```




Figure 5. Finding New Relations.

Every local wordnet has to provide a web service API with the following methods:

1. GetWeightedSynsetsByIli(ILId, ILIversion)
2. GetSynsetById(synsetID)
3. GetSynsetsByLemma(lemma)

¹ For example, the Chinese and the Italian wordnets considered as our case-study use respectively versions 1.6 and 1.5.

² It should be noted, however, that the conversion between different WN versions could not be accurate so the mapping is always proposed with a probability score.

The returned synsets of each method must be formatted in XML following the schema depicted in Figure 6:

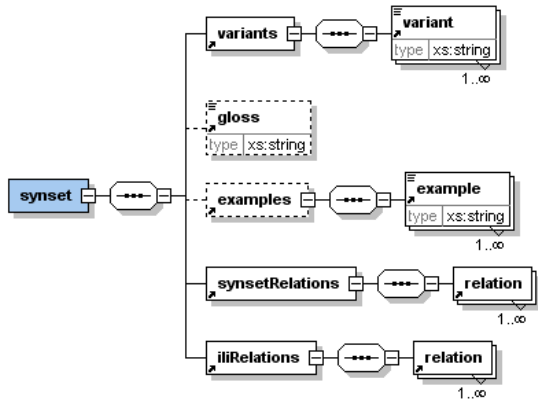


Figure 6. Schema of Wordnet Synsets Returned by WN Web Services.

The scores returned by the method “Get-WeightedSynsetsByIli” are used by our module to calculate the reliability rating for each new proposed relation.

3.3 A Case Study: Cross-fertilization between Italian and Chinese Wordnets.

We explore this idea with a case-study involving the Italian WordNet (Roventini et al., 2003) and the Academia Sinica Bilingual Ontological Wordnet (Sinica BOW, Huang et al., 2004).

The BOW integrates three resources: WordNet, English-Chinese Translation Equivalents Database (ECTED), and SUMO (Suggested Upper Merged Ontology). With the integration of these three key resources, Sinica BOW functions both as an English-Chinese bilingual wordnet and a bilingual lexical access to SUMO. Sinica BOW currently has two bilingual versions, corresponding to WordNet 1.6. and 1.7. Based on these bootstrapped versions, a Chinese Wordnet (CWN, Huang et al. 2005) is under construction with handcrafted senses and lexical semantic relations. For the current experiment, we have used the version linking to WordNet 1.6.

ItalWordNet was realized as an extension of the Italian component of EuroWordNet. It comprises a general component consisting of about 50,000 synsets and terminological wordnets linked to the generic wordnet by means of a specific set of relations. Each synset of ItalWordNet is linked to the Interlingual-Index (ILI).

The two lexicons refer to different versions of the ILI (1.5 for IWN and 1.6 for BOW), thus making it necessary to provide a mapping be-

tween the two versions. On the other hand, no mapping is necessary for the set of synset relations used, since both of them adopt the same set.

For the purposes of evaluating the cross-lingual module, we have developed two web-services for managing a subset of the two resources.

The following Figure shows a very simple example where our procedure discovers and proposes a new meronymy relation for the Italian synset {passaggio, strada, via}. This synset is equivalent to the ILI “road, route” that is ILI-connected with BOW synset “道路, 道, 路” (dao_lu, dao, lu) (Figure 7, A). The Chinese synset has a meronymy relation with the synset “十字路口” (wan) (B). This last synset is equivalent to the ILI “bend, crook, turn” that is ILI-connected with Italian WordNet synset “curvatura, svolta, curva” (C). Therefore the procedure will propose a new candidate meronymy relation between the two Italian WordNet synsets (D).

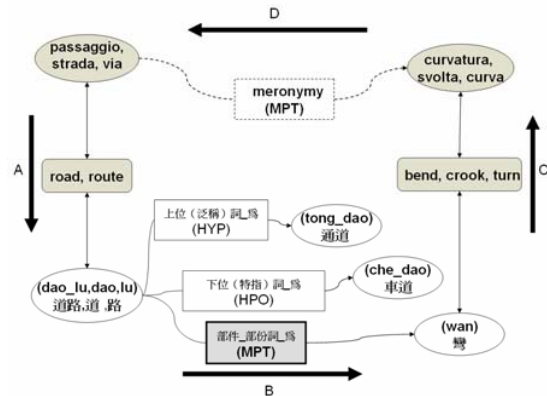


Figure 7. Example of a New Proposed Meronymy Relation for Italian.

3.4 Considerations and Lessons Learned

Given the diversity of the languages for which wordnets exist, we note that it is difficult to implement an operational standard across all typologically different languages. Work on enriching and merging multilingual resources presupposes that the resources involved are all encoded with the same standard. However, even with the best efforts of the NLP community, there are only a small number of language resources encoded in any given standard. In the current work, we presuppose a de-facto standard, i.e. a shared and conventionalized architecture, the WordNet one. Since the WordNet framework is both conventionalized and widely followed, our system is

able to rely on it without resorting to a more substantial and comprehensive standard. In the case, for instance, of integration of lexicons with different underlying linguistic models, the availability of the MILE (Calzolari et al., 2003) was an essential prerequisite of our work. Nevertheless, even from the perspective of the same model, a certain degree of standardization is required, at least at the format level.

From a more general point of view, and even from the perspective of a limited experiment such as the one described in this paper, we must note that the realization of the new vision of distributed and interoperable language resources is strictly intertwined with at least two prerequisites. On the one side, the language resources need to be available over the web; on the other, the language resource community will have to reconsider current distribution policies, and to investigate the possibility of developing an “Open Source” concept for LRs.

4 Conclusion

Our proposal to make distributed wordnets interoperable has the following applications in processing of lexical resources:

- Enriching existing resources: information is often not complete in any given wordnet: by making two wordnets interoperable, we can bootstrap semantic relations and other information from other wordnets.
- Creation of new resources: multilingual lexicons can be bootstrapped by linking different language wordnets through ILI.
- Validation of existing resources: semantic relation information and other synset assignments can be validated when it is reinforced by data from a different wordnet.

In particular, our work can be proposed as a prototype of a web application that would support the Global WordNet Grid initiative (www.globalwordnet.org/gwa/gwa_grid.htm).

Any multilingual process, such as cross-lingual information retrieval, must involve both resources and tools in a specific language and language pairs. For instance, a multilingual query given in Italian but intended for querying English, Chinese, French, German, and Russian texts, can be sent to five different nodes on the Grid for query expansion, as well as performing the query itself. In this way, language specific

query techniques can be applied in parallel to achieve best results that can be integrated in the future. As multilingualism clearly becomes one of the major challenges of the future of web-based knowledge engineering, WordNet emerges as one leading candidate for a shared platform for representing a lexical knowledge model for different languages of the world. This is true even if it has to be recognized that the wordnet model is lacking in some important semantic information (like, for instance, a way to represent the semantic predicate). However, such knowledge and resources are distributed. In order to create a shared multi-lingual knowledge base for cross-lingual processing based on these distributed resources, an initiative to create a grid-like structure has been recently proposed and promoted by the Global WordNet Association, but until now has remained a wishful thinking. The success of this initiative will depend on whether there will be tools to access and manipulate the rich internal semantic structure of distributed multi-lingual WordNets. We believe that our work on LeXFlow offers such a tool to provide inter-operable web-services to access distributed multilingual WordNets on the grid.

This allows us to exploit in a cross-lingual framework the wealth of monolingual lexical information built in the last decade.

5 References

- Nicoletta Calzolari, Francesca Bertagna, Alessandro Lenci and Monica Monachini, editors. 2003. *Standards and Best Practice for Multilingual Computational Lexicons. MILE (the Multilingual ISLE Lexical Entry)*. ISLE CLWG Deliverable D2.2 & 3.2. Pisa.
- Nicoletta Calzolari and Claudia Soria. 2005. A New Paradigm for an Open Distributed Language Resource Infrastructure: the Case of Computational Lexicons. In *Proceedings of the AAAI Spring Symposium “Knowledge Collection from Volunteer Contributors (KCVC05)”*, pages 110-114, Stanford, CA.
- Nicoletta Calzolari. 2006. Technical and Strategic issues on Language Resources for a Research Infrastructure In *Proceedings of the International Symposium on Large-scale Knowledge Resources (LKR2006)*, pages 53-58, Tokyo, Tokyo Institute of Technology.
- Jordi Daudé, Lluís Padró and German Rigau. 2001. A Complete WN1.5 to WN1.6 Mapping. In *Proceedings of NAACL Workshop “WordNet and Other Lexical Resources: Applications, Extensions and*

- Customizations*", pages 83-88, Pittsburg, PA, USA, Association for Computational Linguistics.
- Greg Gulrajani and David Harrison. 2002. SHAWEL: Sharable and Interactive Web-Lexicons. In *Proceedings of the LREC2002 Workshop on Tools and Resources in Field Linguistics*, pages 1-4, Las Palmas, Canary Islands, Spain.
- Chu-Ren Huang, Ru-Yng Chang, and Shiang-Bin Lee. 2004. Sinica BOW (Bilingual Ontological Wordnet): Integration of Bilingual WordNet and SUMO. In *Proceedings of LREC2004*, pages 1553-1556, Lisbon, Portugal.
- Chu-Ren Huang, Chun-Ling Chen, Cui-Xia Weng, Hsiang-Ping Lee, Yong-Xiang Chen and Keh-jiann Chen. 2005. The Sinica Sense Management System: Design and Implementation. *Computational Linguistics and Chinese Language Processing*. 10(4): 417-430.
- Marc Kemps-Snijders, Mark-Jan Nederhof, and Peter Wittenburg. 2006. LEXUS, a web-based tool for manipulating lexical resources. Accepted for publication in *Proceedings of LREC2006*, Genoa, Italy.
- Andrea Marchetti, Maurizio Tesconi, and Salvatore Minutoli. 2005. XFlow: An XML-Based Document-Centric Workflow. In *Proceedings of WISE'05*, pages 290-303, New York, NY, USA.
- Wim Peters, Piek Vossen, Pedro Diez-Orzas, and Geert Adriaens. 1998. Cross-linguistic Alignment of Wordnets with an Inter-Lingual-Index. In Nancy Ide, Daniel Greenstein, and Piek Vossen, editors, Special Issue on EuroWordNet, *Computers and the Humanities*, 32(2-3): 221-251.
- Laurent Romary, Gil Francopoulo, Monica Monachini, and Susanne Salmon-Alt 2006. Lexical Markup Framework (LMF): working to reach a consensual ISO standard on lexicons. Accepted for publication in *Proceedings of LREC2006*, Genoa, Italy.
- Adriana Roventini, Antonietta Alonge, Francesca Bertagna, Nicoletta Calzolari, Christian Girardi, Bernardo Magnini, Rita Marinelli, and Antonio Zampolli. 2003. ItalWordNet: Building a Large Semantic Database for the Automatic Treatment of Italian. In Antonio Zampolli, Nicoletta Calzolari, and Laura Cignoni, editors, *Computational Linguistics in Pisa*, IEPI, Pisa-Roma, pages 745-791.
- Nilda Ruimy, Monica Monachini, Elisabetta Gola, Nicoletta Calzolari, Cristina Del Fiorentino, Marisa Ulivieri, and Sergio Rossi. 2003. A Computational Semantic Lexicon of Italian: SIMPLE. In Antonio Zampolli, Nicoletta Calzolari, and Laura Cignoni, editors, *Computational Linguistics in Pisa*, IEPI, Pisa-Roma, pages 821-864.
- Claudia Soria, Maurizio Tesconi, Francesca Bertagna, Nicoletta Calzolari, Andrea Marchetti, and Monica Monachini. 2006. Moving to Dynamic Computational Lexicons with LeXFlow. Accepted for publication in *Proceedings of LREC2006*, Genoa, Italy.
- Piek Vossen. 1998. Introduction to EuroWordNet. In Nancy Ide, Daniel Greenstein, and Piek Vossen, editors, Special Issue on EuroWordNet, *Computers and the Humanities*, 32(2-3): 73-89.