Intensive Use of Lexicon and Corpus for WSD

Iulia Nica*, M. Antònia Martí
CLiC-Centre de Llenguatge i Computació
University of Barcelona, Spain
* Departament of General Linguistics
University of Iasi, Romania
iulia@clic.fil.ub.es, amarti@ub.edu

Andrés Montoyo, Sonia Vázquez
Dept. of Information Systems and Languages
University of Alicante , Spain
montoyo@dlsi.ua.es, svazquez@dlsi.ua.es

Resumen: El artículo trata sobre el uso de información lingüística en la Desambiguación Semántica Automática (DSA). Proponemos un método de DSA basado en conocimiento y no supervisado, que requiere sólo un corpus amplio, previamente etiquetado a nivel morfológico, y muy poco conocimiento gramatical. El proceso de DSA se realiza a través de los patrones sintácticos en los que una ocurrencia ambigua aparece, en base a la hipótesis de "*almost one sense per syntactic pattern*". Esta integración nos permite extraer información paradigmática y sintagmática del corpus relacionada con la ocurrencia ambigua. Usamos variantes de la información de EuroWordNet asociada a los sentidos y dos algoritmos de DSA. Presentamos los resultados obtenidos en la aplicación del método sobre la tarea *Spanish lexical sample* de Senseval-2. La metodología es fácilmente transferible a otras lenguas.
Palabras clave: Desambiguación Semántica Automática, EuroWordNet

Abstract: The paper addresses the issue of how to use linguistic information in Word Sense Disambiguation (WSD). We introduce a knowledge-driven and unsupervised WSD method that requires only a large corpus previously tagged with POS and very little grammatical knowledge. The WSD process is performed taking into account the syntactic patterns in which the ambiguous occurrence appears, relaying in the hypothesis of "almost one sense per syntactic pattern". This integration allows us to obtain, from corpora, paradigmatic and syntagmatic information related to the ambiguous occurrence. We also use variants of EWN information for word senses and different WSD algorithms. We report the results obtained when applying the method on the Spanish lexical sample task in Senseval-2. This methodology is easily transportable to other languages.
Keywords: Word Sense Disambiguation, EuroWordNet

1 Introduction

The lack of satisfactory results in Word Sense Disambiguation (WSD) claims a deep exam of the process, in order to explore new ways to develop it. Linguistic knowledge could play an important role in this attempt, as empirical evidence shows the greater contribution of information with respect to algorithms in the process of WSD (Pedersen, 2002).

The present paper proposes concrete modalities to strengthen the presence of the linguistic knowledge in WSD, by means of varied exploitation of the corpus and of the lexicon.

We introduce a WSD method based on a fundamental property of natural language, the interaction between the syntagmatic and paradigmatic axes. The basic idea is the development of the WSD process for an ambiguous occurrence integrated in its local syntactic patterns. Our WSD method is knowledge-driven and unsupervised: it uses information from EuroWordNet (EWN, Vossen, 1998) and it does not need a sense-tagged training corpus. It only requires a large corpus, with a minimal preprocessing phase (POS-tagging), and very little grammatical knowledge. Up to now, we have applied the method on Spanish, for nouns disambiguation.

The paper is organised as follows: the previous related work (section 2), our approach to WSD (section 3), the experimentation and its analysis (section 4), finally the conclusions and future work (section 5).

2 *Previous Related Work*

In the bibliography, there are WSD methods that meet ours from some point of view.

The syntactic information exploited for WSD has been limited generally to verb-subject and verb-object relations (Ng, 1996; Martínez *et al*., 2002, etc.), with few exceptions (Lin, 1997; Stetina *et al*., 1998). The relations were acquired from syntactically and semantically annotated corpora.

The use of functional words contiguous to the ambiguous occurrence have been done especially in the example-based approach, and so it has been related to a sense tagged corpus (Pedersen, 2001; Yarowsky and Florian, 2002; Mihalcea, 2002; Hoste *et al*., 2002, etc.).

There are some proposals to obtain sense-tagged examples from corpora or Internet by replacing a sense-tagged word with monosemous words related to it along lexical-relations in WordNet (WN) (Leacock *et al.,* 1998; Mihalcea and Moldovan, 1999). They are limited to the existence of such monosemous related words.

Our work is closer to corpus-based methods of Montemagni *et al*. (1996), Federici *et al*. (2000), defined as "Paradigm-driven Approach" to WSD, and of Agirre and Martínez (2001). In these methods there are combined paradigmatic variants for the two lexical content positions of what we call syntactic pattern. The combination is performed only for verb-argument relations and on syntactic patterns already tagged, at syntactic and sense levels.

Our proposal is independent on a syntactically and semantically tagged corpus, and it uses different syntactic relations involving nouns. The method works with real examples in texts, from which it obtains word sets related to the focused word into the syntactic pattern. Furthermore, the obtained sense-tagged patterns can be reused in future disambiguation tasks; as it works in a good percentage on the local context, with several syntactic patterns, our method limits the data-sparseness problem that affects the methods which consider the whole sentence.

3 *Approach to WSD*

One of the critical points in WSD is the lack of information to perform the process. In the case of knowledge-based approach, the lack derives from the gap between lexicon and corpus (Kilgarriff, 1998): paradigmatic information in the lexicons for word senses *vs.* syntagmatic information from the context of the ambiguous occurrence. Instead the corpus-based systems need a huge amount of sense-tagged examples for the word to be disambiguated: the knowledge acquisition bottleneck and data sparseness problems.

We address the issue of information in WSD from the knowledge-based approach, with the purpose of identifying linguistic information useful for the WSD process and how to better exploit the information there involved. Words usage is registered into corpora, thus corpora are a valuable source of linguistic knowledge, insufficiently exploited for WSD tasks.

The WSD process consists in a mapping between two kinds of information: the information associated to the ambiguous occurrence in the corpus and the information related to word senses in the lexical source. Thus, in order to improve the level of the WSD process, our basic strategy is to increase the probabilities of these two sets of information to meet each other. We investigate three possible directions to do it: 1) the enrichment of the information associated to the ambiguous occurrence, by acquiring especially paradigmatic information for it (section 3.1.); 2) the refinement of the information related to word senses provided in the lexicon (section 3.2.); 3) the use of different modalities to map the two sets of information (section 3.3.).

3.1 Rich Information for the Ambiguous Occurrence

Along the first direction, we investigate the possibility to reverse the disambiguation process: operate on paradigmatic instead of syntagmatic information. With this purpose, we extract paradigmatic information associated to a given ambiguous word and then we map it to the paradigmatic information from the lexicon for the word senses. This mapping is performed by a set of heuristics (section 3.3.). The paradigmatic information related to the ambiguous occurrence is obtained by exploiting the interaction that holds in natural language between the syntagmatic and paradigmatic axes: semantically similar words can substitute each other in the same context and, inversely, words that can commute in a context have a good probability to be semantically close.

The starting syntagmatic data in this operation leads us to the issue of local context. For the formal treatment of the context, we introduce the term of syntactic pattern: a triplet X-R-Y, formed by two lexical content units X and Y (nouns, adjectives, verbs, adverbs) and a relational element R (R=∅ inclusively), which corresponds to a syntactic relation between X and Y.

In order to identify occurrences for a noun X in EWN and for their syntactic patterns, we work on a POS-tagged corpus that we call "search corpus". The identification of the syntactic patterns is done following criteria of structure and of frequency. We predefine a list of basic patterns: [N ADJ], [ADJ N], [N PART], [PART N], [N CONJ N], [N PREP N], [N, N], and also search schemes in order to identify discontinuous realisations of patterns inside complex sequences. We introduce some frequency filters on the patterns we obtain: on the one hand, we impose the condition on the potential patterns to repeat into the corpus; on the other hand, we eliminate the ones with more than 1000 substitutes for the word to be disambiguated. In this way, we obtain syntactic patterns P_k for the noun X.

The integration of the ambiguous occurrence in a syntactic pattern allows us to identify, into the corpus, information of paradigmatic and syntagmatic type associated to the occurrence inside the pattern: the sets S1 and S2 below.

-**S_1** is the set of nouns that can substitute X into the considered syntactic pattern. We obtain them by fixing the syntactic pattern at lemma and morphosyntactic levels, and letting variable only the position of X at lemma level. From these substitutes, we keep the more frequent 20 inside the pattern.

-**S_2** is the set of the cooccurring nouns with the syntactic pattern into a sentence. From these nouns, we keep the first 10 more frequently cooccurring with the pattern.

We obtain thus the sets S_{1k} and S_{2k} corresponding to X inside the syntactic patterns P_k.

At the basis of this approach it lays the hypothesis that the sense of an ambiguous occurrence is fundamentally determined by the words from the context with which it establishes syntactic relations. Thus the integration of the occurrence into its syntactic pattern is a first approximation to its sense by reducing the word polysemy: we consider there is a tendency of words to be monosemous inside a syntactic pattern (the "quasi one sense per syntactic pattern" claim).

We also use the information from the sentential context of the occurrence to be disambiguated: the sets of nouns of the sentence, that we call S_3.

The final information we collect for the ambiguous occurrence X is the one extracted from corpus by using all its syntactic patterns and the one provided by the sentential context: sets S_{1k} and S_{2k} (for every k such that exists a syntactic pattern P_k containing X), respectively S_3.

3.2 Refining the Characterisation for Word Senses

We have also investigated how to enrich and diversify the information associated to word senses in the lexicon: we analysed how to design different sense characterisations of EWN synsets by means of the information contained there.

We derive an adaptation of the lexicon in the following way: for every sense X_i of a given word X in EWN, we extract the set of nouns related to it in EWN along the different lexical-semantic relations; we will note the obtained set with SD_i. That is, for every sense X_i of X, we obtain the synsets related to it by hyperonymy, hyponymy, holonymy, etc., until using all the lexical-semantic relations in EWN. For each sense we group in a set all the synsets obtained and we eliminate the common elements (at lemma level), obtaining so the disjunctive sets SD_i. As the elements of the set SD_i are related exclusively with the sense X_i, they become sense discriminators for X_i with respect to EWN. For this reason, we called the obtained lexical device "Sense Discriminators" (SD).

We have thus two characterisations for word senses based on EWN: the classical one, of hierarchical type, based on the position of word senses into the IS-A taxonomy, and the new one, Sense Discriminators, of flat structure, which involves all the lexical-semantic relations in EWN.

3.3 Two WSD Algorithms

For the mapping between the information associated to word senses and the one related to the ambiguous occurrence, we use two different WSD algorithms that exploit the variants of EWN. We present them below.

A_1: The Specificity Mark algorithm (Montoyo and Palomar, 2000). It works on the original form of EWN. The intuitive base of this algorithm is: the more common information two concepts share the more related they will be. In EWN, the common information shared by two concepts corresponds to the father concept of both in the hierarchy, called Specificity Mark (SM) by the authors. The heuristic takes as input a noun set and looks for the SM in EuroWordNet with the bigger density of input words in its subtree. It chooses as correct for every input word the sense situated in the sub-tree of the SM so identified, and it lets undisambiguated the words without senses in this subtree. In its initial version, the algorithm of Specificity Marks was applied on the nouns contained in the sentence of the ambiguous occurrence. We generalise it here to be used on any set of input nouns.

A_2: The Commutative Test algorithm (Nica *et al.*, 2003). It is related to the Sense Discriminators device. At the basis of the algorithm it lays the hypothesis that if two words can commute in a given context, they have a good probability to be semantically close. In terms of our adaptation of EWN and of our approximation to local context, this means that if an ambiguous occurrence can be substituted in its syntactic patterns by a sense discriminator, then it can have the sense associated to that sense discriminator. We call this algorithm the Commutative Test (CT). In order to reduce the computational cost of this substitution operation, we perform an equivalent process: We previously extract, from corpus, the possible substitutes of the ambiguous occurrence in a syntagmatic pattern, and then we intersect this set with every set of sense discriminators; the senses for which the intersection is not empty can be assigned to the occurrence. We use here a generalisation of the CT, to be applied on any set of words related to the ambiguous occurrence.

3.4 WSD System

The method runs through the following steps:
1) the identification of the syntactic patterns for the ambiguous occurrence;
2) the extraction, from corpus, of information related to the ambiguous occurrence;
3) the application of the WSD algorithms on the information previously obtained;
4) the final decision on sense assignment.

We use a WSD system that incorporates several heuristics as voters. The WSD heuristics are determined by the combination of a set S_i associated to the ambiguous occurrence (from $\mathbf{S_1}$, $\mathbf{S_2}$, $\mathbf{S_3}$ in section 2.1) and an algorithm A_j (from $\mathbf{A_1}$, $\mathbf{A_2}$, in section 2.3).

We divide the heuristics in two groups:

-**Heuristics I** are the four heuristics based on the syntactic patterns: $H_{1k} = (S_{1k}, A_1)$, $H_{2k} = (S_{2k}, A_1)$, $H_{3k} = (S_{1k}, A_2)$, $H_{4k} = (S_{2k}, A_2)$.

-**Heuristics II** are the heuristics based on the sentence: $H_1 = (S_3, A_1)$, $H_2 = (S_3, A_2)$.

The modality for the combination of the heuristics into the WSD system has been established on the basis of their individual evaluation (section 4).

For every pattern, we implement the four heuristics of group I and choose the most voted sense (with 2/2, 3/3, 4/4 or 3/4 of responses) or, if any sense satisfy this condition, we choose the sense proposed by the most precise heuristic. We then intersect the sense proposed by the different patterns.

Turning to the heuristics related to the sentence, if we have answers from both heuristics, we do the intersection between the proposed senses. If the intersection is null, we take the answer of H_1.

For the final sense assignment, we first apply the heuristics from group I and after the ones from group II. Otherwise, if there are no identified patterns or if there are no answers form the patterns, we take the proposals of group II. If there is no answer from H_2, we take the one from H_1. In case of parity between two or more senses, we choose the proposal of the most precise heuristic (of all six) between those with a response.

3.5 Example

We illustrate the method for noun *órgano* in the following example from Senseval-2 (occurrence number 75):

> *Un informe del <head>órgano</head> de gobierno de los jueces advierte de que no pretende una modificación legislativa, sino proponer soluciones a problemas del nuevo Código, pues, señala, "provoca la comisión de delitos, tiene penas desproporcionadas y ha generado una situación penitenciaria que alcanzará cotas insostenibles".*

The steps of the disambiguation process are:

PREPROCESSING:

a.Extraction of Sense Discriminators sets

In EWN, *órgano* has five senses[1]:

órgano_1: 'part of a plant';
órgano_2: 'governmental agency, instrument';
órgano_3: 'functional part of an animal';
órgano_4: 'musical instrument'
órgano_5: 'newspaper'.

Correspondingly, we obtain from the EWN hierarchy the following Sense Discriminators sets:

SD_1: {*órgano vegetal, flor, pera, semilla,…*}
SD_2: {*agencia, unidad administrativa, ...*}
SD_3: {*parte del cuerpo, trozo, ojo, …*}
SD_4: {*instrumento de viento, teclado, …*}
SD_5: {*periódico, publicación, número, …*}

b.Input text POS-tagging

STEP 1. Syntactic patterns identification of the ambiguous occurrence

We particularize for *órgano* the general predefined basic patterns and search schemes.

1a. Using these particularised search schemes, we find the following sequence: [*informe*-N *de*-PREP *órgano*-N *de*-PREP *gobierno*-N].

1b. From here, we extract two basic patterns: P1=[*informe*-N *de*-PREP *órgano*-N] and P2=[*órgano*-N *de*-PREP *gobierno*-N].

STEP 2. Extraction of information associated to the ambiguous occurrence

2a. Extraction of paradigmatic information from corpus

In order to extract the paradigm corresponding to the position of *órgano* in each of the two syntactic patterns previously identified, we let vary, at lemma level, the position of *órgano* in the two patterns P_1 and P_2 respectively. With the help of the particularised search schemes, we then look in the corpus for the possible nouns as X in any of the possible realisations of these two patterns. We obtain two sets, corresponding to P_1 and P_2:

S_{11}: {*gestión, comisión, policía, prensa, servicio, organización, experto, coyuntura, organismo, ponencia, autoridad, auditoría, agencia, perito, intervención, observador, conclusión, situación, fiscalización, grupo, emisora, candidatura*, ...}

S_{21}: {*jefe, programa, año, órgano, formación, equipo, partido, coalición, representante, acción, miembro, cambio, comité, parte, alianza, período, pacto, funcionario, crisis, responsabilidad, mes, alternativa, acuerdo*, ...}

2b. Extraction of syntagmatic information from corpus

We look into the corpus for the sentences with the pattern P_1 and, separately, for the sentences with the pattern P_2. The two sets corresponding to P_1 and P_2 are the following:

S_{12}: {*tráfico, sugerencia, producto, observación, mención, medida, justicia, estupefaciente, estilo, desvío, consumo, Junta, anteproyecto, Departamento_de_Justicia, ...*}

S_{22} ={*juez, magistrado, presidente, CGPJ, EFE, fuente, acuerdo, miembro, reunión, poder, comisión, ciudad, Senado, representante, vicepresidente, texto, función, forma,...*}

2c. Extraction of syntagmatic information from the sentential context:

The nouns of the sentence are:

S_3: {*órgano, informe, gobierno, juez, modificación, solución, código, comisión, delito, penas, situación, cota*}

STEP 3. Application of the WSD algorithms on the information associated to the ambiguous occurrence

In table 1, we synthesise the proposals for the sense of *órgano* from all the heuristics:

P_1				P_2				Sentence	
H_{11}	H_{21}	H_{31}	H_{41}	H_{12}	H_{22}	H_{32}	H_{42}	H_1	H_2
s4 s5	-	s2	s2	s2	-	s2	-	s1 s2	s2

Table 1: Results for the occurrence *órgano*#75

[1] The pseudo-definitions are ours.

STEP 4. Final sense assignment

In this case, we obtain the sense 2 from pattern P_2 and from the sentence, so we assign sense 2 from EWN to the occurrence of *órgano*, which corresponds to sense 3 in the Senseval-2 dictionary.

4 Experiments and Discussion

We have applied our method on the nouns in the test corpus from the Spanish Senseval-2 exercise, in order to obtain a more objective evaluation. We present here the results obtained when using as search corpus EFE and LEXESP (over 75 million words) and the POS-tagger (Civit, 2003).

In the experimentation, we used the 1.6 variant of EWN, whose mapping to the reference dictionary from Senseval-2 (Rigau *et al.*, 2001) is partial: for the 17 test words, the mapping covers 78,31% of the senses in EWN and 88,67% of the senses in the dictionary. Thus for only 688 from the 799 nominal test occurrences (that is 86,10%) we could verify the sense assignation in terms of EWN. In the experimentation, we have continued only with the 688 occurrences that do have a sense assignation in terms of our EWN variant. We do this in order to have a reliable evaluation of our method, independently on the mapping limitation.

For the 688 occurrences in the Senseval-2 test corpus, we have obtained 803 syntactic patterns and, from these, 318 filtered patterns corresponding to 267 occurrences. That is we have coverage with patterns of 54,71% on the 688 analysed occurrences. The evaluation of the individual heuristics is presented in table2, in terms of the syntactic patterns identified for the occurrences.

Heuristics		Precision	Coverage (patterns)	Coverage (occurrences)
Heur. I	H_{1k}	35,45%	94,02%	43,45%
	H_{2k}	32,14%	52,83%	24,41%
	H_{3k}	58,92%	52,83%	24,41%
	H_{4k}	68,33%	37,73%	17,44%
Heur. II	H_1	28,93%	-	96,94%
	H_2	62,74%	-	6,38%

Table 2: Evaluation of the individual heuristics

The final results when combining the heuristics as indicated in section 3.4 are presented in table 3.

Heuristics	Precision	Recall	Coverage
Heuristics I	92,59%	3,63%	3,92%
Heuristics II	31,63%	30,66%	96,94%
Heuristics I + II	33,28%	32,26%	96,94%

Table 3: Final results

The evaluation in table 4 (Rigau *et al.*, 2001) indicates a low level of performance of our method with respect to the level reached for nouns in Senseval-2. From the systems in Senseval-2, we give a special attention to the knowledge-based system of the University of Alicante, constructed with the Specificity Marks algorithm applied on the sentence (here heuristic H_1) complemented by seven heuristics.

WSD system			Precision	Recall
Our system			33,28%	32,26%
Senseval-2	Corpus-based	Superior	66%	66%
		Inferior	55%	55%
	Knowledge-based (UA)		55%	43,5%

Table 4: Comparison with Senseval-2 (nouns)

The low results are related to the limited participation of the patterns heuristics (heuristics I) to the disambiguation process (table3). This limitation is principally due to the following facts: we have not used patterns covering argumental and syntactic relations of nouns with verbs, nor some frequent search schemes (because their decomposition into basic patterns is not trivial and needs a special study); we filtered too strictly the patterns by imposing them to have less than 1000 substitutes for the focused word (the number of patterns dramatically reduced from 803 to 318, that is at 39, 60%); the insufficient dimension of the corpus we used (75 million words) negatively affected both coverage and precision.

Our principal purpose in this experimentation has been to test different aspects of our proposal for WSD:

a) the validity of the hypothesis “quasi one sense for syntactic pattern”. We have tested it on the iterative syntactic patterns (53) in the Senseval-2 test corpus. Even data is very limited, it seems that there is a tendency of the syntactic patterns to associate with a unique word sense: 49 cases on 53 (92,4% of the patterns); in the other 4 cases, the word ambiguity inside the syntactic patterns reduces to two senses. This is a very partial confirmation of our strategy to integrate the ambiguous occurrences into syntactic patterns as a first step towards their disambiguation.

b) the usefulness of the information obtained from corpus for the ambiguous occurrence starting from its syntactic pattern. The results indicate that there can be done WSD using only syntactic patterns and that the use of syntactic patterns improves the WSD level. At the same time, the information extracted form corpus, associated to the ambiguous occurrence, improves the performance of the WSD algorithms on the occurrence sentence.

c) the usefulness of the Sense Discriminators device and of the associated algorithm, the Commutative Test. The tests have shown that we do can make WSD by means of the Sense Discriminators device and the Commutative Test, with superior precision even if with lower coverage (table 2).

5 Conclusions and Future Work

We present in this article a method for Word Sense Disambiguation that takes as unit to disambiguate an ambiguous occurrence integrated into a syntagmatic pattern. It starts from the hypothesis that this integration decisively reduces the polysemy of the ambiguous occurrence: the “quasi one sense per syntactic pattern”. The strategy exploits the semantic interaction between the paradigmatic and syntagmatic axes which holds in natural language. This property allows us to identify in the corpus information related to the occurrence to be disambiguated, and apply on it the information from the lexicon (EWN). At the same time, we use different word senses characterisations starting from EWN and different WSD algorithms to exploit them.

An important characteristic of our proposal is the independence on a corpus tagged at sense or syntactic level. The method only needs a POS-tagger, for the analysis both of the sentence in which occurs the token to be disambiguated and of the search corpus, so it can easily be adapted to other languages. Furthermore, it has an expansive potential, as it allows the simultaneous (pre-)disambiguation of different occurrences of a word in any text, when they appear in the same pattern.

We performed a testing of our method in the conditions of the Spanish Senseval-2 exercise, with the following results: 33,28% precision, 32,26% recall, 96,94% coverage.

As future work, we have to overcome the present limitations of the method with respect both to coverage and precision. In order to improve the coverage, we first have to consider patterns with verbs and relax the filters on the patterns. For the improvement of precision, we are currently comparing different combinations of the heuristics. At the same time, we are investigating the syntactic pattern acquisition from untagged corpora previously to the WSD process and further, with the help of these patterns, the acquisition of sense clues. We are also analysing the search by changing the corpus with Internet.

References

Agirre, E. and D. Martínez. 2001. Learning class-to-class selectional preferences. *Proceedings of the ACL* CONLL’2001 Workshop, Tolouse

Civit, M. 2003. *Criterios de etiquetación y desambiguación morfosintáctica de corpus en español*, PH.D.Thesis, Universidad de Barcelona

Federici, S., S. Montemagni and V. Pirelli. 2000. ROMANSEVAL: Results for Italian by SENSE. *Computers and the Humanities. Special Issue: Evaluating WSD Programs, 34 (1-2)*

Hoste, V., I. Hendrickx, W. Daelemans and A. van den Bosch, 2002. Parameter optimisation for machine-learning of WSD. *Natural Language Engineering, 8 (4)*

Kilgariff, A. 1998. Bridging the gap between lexicon and corpus: convergence of formalisms. *Proceedings of LREC'1998*, Granada

Leacock, C., M. Chodorow and G.A. Miler. 1998. Using Corpus Statistics and WordNet Relations for Sense Identification. *Computational Linguistics. Special Issue on Word Sense Disambiguation, 24(1)*

Lin, D., 1997. Using Syntactic Dependency as Local *Context to* Resolve Word Sense Ambiguity. *Proceedings of ACL and EACL'97*, Morgan Kaufman Publishers, San Francisco

Martínez, D., E. Agirre and L. Màrquez. 2002. Syntactic Features for High Precision Word Sense Disambiguation. *Proceedings of the 19th International Conference on Computational Linguistics (COLING 2002),* Taipei

Mihalcea, R. 2002. WSD with pattern learning and feature selection. *Natural Language Engineering, 8(4),* Cambridge University Press

Mihalcea, R. and D. Moldovan. 1999. An Automatic Method for Generating Sense Tagged Corpora. *Proceedings of AAAI '99*, Orlando

Montemagni, S., S. Federici and V. Pirelli. 1996. Example-based Word Sense Disambiguation: a Paradigm-driven Approach. *Proceedings of EURALEX '96*, Göteborg

Montoyo, A. and M. Palomar. 2000. Word Sense Disambiguation with Specification Marks in Unrestricted Texts. *Proceedings of 11th International Workshop on DEXA*, Greenwich

Ng, H.T. and H.B. Lee. 1996. Integrating Multiple Knowledge Sources to Disambiguate Word Sense: An Exemplar-Based Approach. *Proceedings of the 34th Annual Meeting of the ACL*

Nica, I., M. A. Martí and A. Montoyo, 2003. Colaboración entre información paradigmática y sintagmática en la Desambiguación Semántica Automática. *Procesamiento del Lenguaje Natural, 31*

Pedersen, T. 2001. A decision tree of bigrams is an accurate predictor of word sense. *Proceedings of NAACL 2001*, Pittsburg

Pedersen, T. 2002. A Baseline Methodology for Word Sense Disambiguation. *Proceedings of the Third International Conference on Intelligent Text Processing and Computational Linguistics*, Mexico City

Rigau, G., M. Taulé, A. Fernández and J. Gonzalo. 2001. Framework and Results for the Spanish SENSEVAL. Preiss, J. and D. Yarowsky (eds.), *Proceedings of the SENSEVAL-2 Workshop.* ACL 2001/EACL 2001, Toulouse

Stetina, J., S. Kurohashi and M. Nagao. 1998. General WSD Method Based on a Full Sentential Context. *Proceedings of COLING-ACL Workshop*, Montreal

Vossen, P., 1998 (ed.). *EUROWORDNET. A Multilingual Database with Lexical Semantic Networks*, Kluwer Academic Publishers, Dordrecht

Yarowsky, D. and R. Florian. 2002. Evaluating sense disambiguation across diverse parameter spaces. *Natural Language Engineering, 8(4),* Cambridge University Press