Identification of Hyponyms, Hyperonyms, Meronyms and Holonyms in Medical Texts: a Cognitive Approach

César Aguilar and Olga Acosta
Pontificia Universidad Católica de Chile, Santiago de Chile
{caguilara, oacostal}@uc.cl

Abstract

In this paper we propose a method for extracting hyponymy/hyperonymy and meronymy/holonymy relations from a medical corpus in Spanish. Our method considers the following steps: (a) the recognition of analytical *definitional contexts* in our corpus, (b) the identification of noun phrases (NP) that can represent hyponymy-hyperonymy or meronymy-holonymy relations taking into consideration only hyperonyms modified by relational adjectives. We suppose that relational adjectives introduce good candidates to hyponyms derived from a hyperonym. On the other hand, locative adverbs are useful elements for recognizing good candidates to concrete entities which can be in a meronymy-holonymy relation when are in a head position of phrase and modified by a relational adjective.

1 Introduction

The identification of lexical-semantic relations expressed in texts is the main goal of much of the research to date in NLP, particularly those oriented to the building of ontologies and taxonomies. A paradigmatic example of this kind of research is the volume prepared by Buitelaar, Cimiano and Magnini [2005]. These authors develop a complete methodology for building ontologies based on the extraction of conceptual information from text corpora. These authors conceive a layer cake scheme that represents six concrete tasks which should be covered for the building of an ontology:

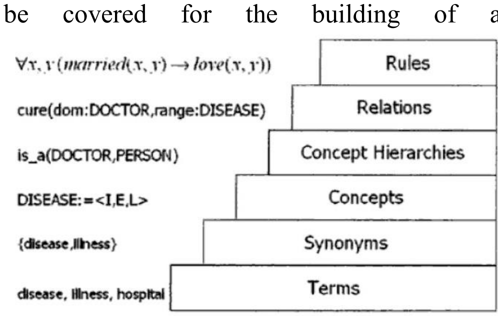

Figure 1. Ontology learning layer cake according to Buitelaar, Cimiano and Magnani [2005: 5]

According to this scheme, we can observe an increase on the degree of complexity implied in each step required for building an ontology. Thus, the first task is the identification of a set of specific terms to a certain knowledge domain (in this case, a medical domain). Then, synonyms are identified (e.g. disease/illness). Given both sets of terms and synonyms, the following task is to recognize concepts expressed in specific text fragments. From these concepts, lexical relations can be inferred in order to configure a conceptual hierarchy.

In the immediate levels, the authors consider the description of lexical-semantic relations configured according to models as the WordNet lexicon [Fellbaum, 1998], or the generative lexicon of Pustejovsky [1995], among others. Finally, the top level represents the formulation of axioms able to infer universal rules that underlie in a conceptualization process. In words of the authors, to achieve this level is the final step for building a true ontology.

In line with Buitelaar, Cimiano and Magnini, researchers as Tsujii and Ananiadou [2005], or Ceusters and Smith [2010] explain the importance of terminological extraction as a previous phase for building ontologies, mainly in a medical domain. In fact, according to Smith [2004], an ontology can be organized around two lexical-semantic relations: hyponymy/hyperonymy and meronymy/holonymy relations. Smith argues that these relations, from a philosophical point of view, are basic and universal.

According to these ideas, in this paper we describe a general methodology for recognizing terms in medical texts in Spanish where hyponymy/hyperonymy and meronymy/holonymy relations can be found. We organize the explanation of our methodology as follows: in the section (2) we offer a state of the art about the automatic extraction of these relations. Then, in the section (3) we explain the theoretical framework considered for this methodology: the prototype theory, the process of categorization, as well as the notions of *spacial scene* and *axial properties* formulated by cognitive sciences [Croft and Cruse 2004; Evans and Green 2006; Evans, 2007]. Immediately, in section (4) we summarize our method describing the corpus, heuristics and measures employed for detecting noun phrases (NP) that contain relational adjectives and locative adverbs, considering that relational adjectives introduce good candidates to hyponyms derived from a hyperonym, while locative adverbs provide a good mechanism for identifying candidates to concrete entities that can be as noun heads of noun phrases where there is a meronymy-holonymy relation (e.g. mucosa abdominal). Finally, in (5) we expose our preliminary results.

2 Extraction of Lexical-Semantic Relations

For Jurafsky and Martin, the recognition of lexical-semantic relations in corpora is a task derived from the relation detection and classification [2009: 734-742], which is oriented to find and classify semantic relations among the entities discovered in a given text. Two lexical relations widely considered in this kind of extraction are the hyponymy/hyperonymy and meronymy/holonymy.

2.1 Extraction of Hyponymy/Hypernymy

Since the pioneering work of Hearst [1992], or Wilks, Slator and Guthrie [1995], the most considered relation in this kind of extraction is the hyponym/hypernym. In the case of Hearst, her experiment offers a significant method for identifying lexical-syntactic patterns associated to hyponyms in large-corpora. Based on Hearts's experiment, there are other alternative approaches:

- **Clustering:** this approach emphasizes the distribution of context in corpus. According to this approach, words are characterized by its context and grouped by its similarity between contexts. One representative work about this approach is Faure and Nedellec [1998].
- **Finding patterns using the Web:** in this approach new characteristic patterns and instances of the lexical relation of interest are extracted taking into account the Web as a huge source of textual information, according to the experiment performed by Pantel and Pennacchiotti [2006].
- **Machine learning:** Finally, Snow, Jurafsky and Ng [2006] proposed an approach considering the application of machine learning methods, oriented to recognize useful patterns employing dependency paths.

2.2 Extraction of Meronymy/Holonymy

The first attempt for extracting automatically this kind of relation is the work of Berland and Charniak [1999]. They focused on genitive patterns identified through the employ of lexical seeds inserted in a part-whole relation with other words (e.g.: *the **basement** of a **building***). For this search, they use a news corpus of 100,000,000 words, and generate an ordered list of part-whole candidates inferred by a log-likelihood metric. Both researchers obtained a level of accuracy around 55% respect to the words associated in a relation part-whole with such seeds. Once achieved these results, the authors compare them with the meronyms associated with WordNet [Fellbaum 1998] in order to determine precision.

On the other hand, Girju, Badulescu and Moldovan [2006] conceived a different method considering a large list (around 50) of possible patterns. They elaborated a corpora with texts taken from LA Times and Wall Street Journal (WSJ). For detecting these patterns, they designed an algorithm named Iterative Semantic Specialization Learning (ISSL), which introduces a process of machine learning for recognizing new meronymy sequences in corpus, derived from the mentioned patterns. ISSL achieves a level of precision almost 83% in the LA Times, and 79% in the WSJ. In contrast, the level of recall is 79% for the first corpus, and 85% for the second corpus. Finally, they also compare the meronymy relations identified with a set of meronyms from WordNet.

3 Theoretical assumptions

The study of lexical-semantic relations is a classical topic in psycolinguistics, cognitive sciences and artificial intelligence [Croft and Cruse 2004; Evans and Green 2006; Jurafsky and Martin 2009]. In computational linguistics, the impact of this topic is recognizable since the creation and development of WordNet [Fellbaum, 1998]. Nowadays, the use of WordNet as a viable knowledge source for evaluating the results generated from any experiment for relation detection is a common task.

Many experiments in relation detection take advantage of the linguistic knowledge compiled in WordNet, and focus their efforts only in the implementation of hybrid methods — combining the use of linguistic patterns with probabilistics techiques— for identifying lexical units involved on these relations, e.g., the experiment of Hearts [1992]. However, there are few works —e.g., the case of Girju, Badulescu and Moldovan [2006]—, which sketch a theoretical framework based in the study made by Winston, Chaffin and Herrmann [1987], in order to orient their experiment.

We believe that the work of Girju, Badulescu and Moldovan is an important advance in the conception of relation detection tasks, because takes into account theoretical criteria came from detailed studies performed by cognitive scientists about conceptualization in humans [Evans and Green 2006; Evans, 2007]. In our paper, we assume three basic theories conceived within cognitive sciences: (i) the prototype theory, (ii) the process of categorization, and (iii) the notion of *axial properties*.

3.1 Prototype Theory

The prototype theory was proposed mainly by Rosch [1978]. According to her, the instances of a concept differ in the degree to which they share certain properties, and consequently show a variation respect to the degree of representation of such concept. An example is: if we formulate a unitary description of the concept *cup*, this might consist of the following five properties: (1) concrete object, (2) concave, (3) can hold liquids, (4) has a handle and, (5) can be used to drink hot liquid out of [Smith and Medin 1983].

An important question here is: are all the properties true of all the *cups*? We might think properties 1-3 are true for all the *cups*, but 4 and 5 are controversial. Then, if we remove properties 4 and 5 from the description, some non-cups are included, e.g., *bowls*. Such considerations allow to argue that there are a great deal of concepts where is difficult to posit a unitary description. Thus, prototype theory provided a new view in which a unitary description of concepts remains, but the properties are true of most, but not all members.

3.2 Categorization Processes

Rosch [1978] proposes two principles in order to build a system of categories. The first refers to the function of this system, which must provide a maximum of information with the least cognitive effort. The second emphasizes that the perceived world (not-metaphysical) has structure. Maximum information with least cognitive effort is achieved if categories reflect the structure of the perceived world as better as possible. Both the cognitive economy principle and the structure of perceived world have important implications in the construction of a system of categories.

Rosch conceives two dimensions in this system: vertical and horizontal. Vertical dimension refers to the category's level of inclusiveness, that is, the subsumption relation between different categories. In contrast, horizontal dimension focuses on segmentation of categories in the same level of inclusiveness. It is possible to see better how work the vertical and horizontal dimensions in the following figure:

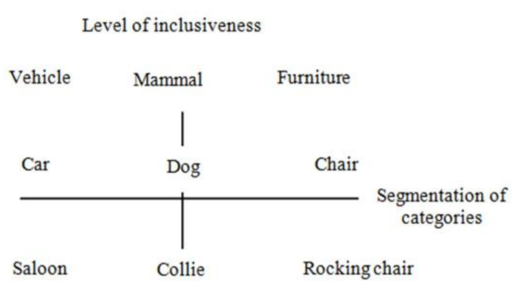

Figure 2. The human categorization system [Evans and Green 2006: 256]

According to the Figure 2, the category *furniture* is more inclusive than the category *chair*, because it includes other entities like, for example, *desk* or *table* in addition to *chair*. In turn, *chair* is more inclusive than *rocking chair* because it includes also other types of chairs. Finally, *rocking chair* as category only includes rocking chairs, and for this reason is the least inclusive level.

The level between the most inclusive and least inclusive level is called the *basic level*, and categories at this level are named *basic-level categories*. Categories higher up the vertical axis, which provide less detail, are called *superordinate categories*. Those lower down the vertical axis, which provide more detail, are identified as *subordinate categories*.

3.3 Spacial Scene and Axial Properties

The conception of space in humans is an important topic in cognitive sciences, and it is referred to as *spatial scene*. According to Evans [2007], a spatial scene is a linguistic unit that contains information about our spatial experience, and it is structured according to four parameters: a *figure* (or *trajector*), a *referent object* (that is, a *landmark*), a *region* and —in certain cases— a *secondary reference object*. These two reference objects configure a *reference frame*.

We can understand better this analyzing the following example: *A car is parked behind the school.* Here, *a car* is the *figure* and *the school* is the *reference object*. Respect to the *region*, this is established by the combination of the preposition which sketches a spatial relation with the reference object. Finally, such relation encodes the location of the *figure*.

Related to the spatial scene, Evans (2007) points out the existence of *axial properties,* that is, a set of spatial features associated with a specific *reference object*. Considering again the sentence *a car is parked near to the school*, we can identify the location of the *car* searching for it in the region near to the *school*. Therefore, this search can be performed because the referent object (*the school*) has a set of *axial properties*: front, back and side areas. These *axial properties* configure all spatial relations.

4 Methodology proposed

We expose here our methodology for extracting candidates of hyponyms, hypernyms, meronyms and holonyms, based in the results obtained by Acosta, Aguilar and Sierra [2013], as well as Acosta and Aguilar [2015].

4.1 Identification of Definitional Contexts

The first step to cover in our methodology is the identification of textual fragments named *analytical definitional contexts* (or DCs), which are constituted by a term, a definition, and linguistic or metalinguistic forms, such as verbal phrases, typographical markers and/or pragmatic patterns [Sierrra *et al*., 2008]. An example of DC is:

> **La energía primaria**, en términos generales, se define como aquel recurso energético que no ha sufrido transformación alguna, con excepción de su extracción. (Eng. The **primary energy**, in general terms, is defined as a resource that has not been affected for any transformation, with the exception of its extraction.)

In this example, the DC is formed by the term *energía primaria* (Eng.: *primary energy*), the definition *aquel recurso…* (Eng.: *that resource that*…) and the predicative phrase *se define como* (Engl.: *is defined as*). In addtition, we can identify other characteristic units such as the pragmatic pattern *en términos generales* (Eng.: *in general terms*) and the typographical marker (bold font) that in this case emphasizes the presence of the term.

Taking into account these elements, the bakcbone of a DC is the sequence conformed by the term, the predicative phrase and the definition. We can observe how this sequence operates in the following example:

[Conjunctivitis [Term]] [is [Predicative Phrase]] [an inflammation of the conjunctiva of the eye [Definition]]

In this example, we observe the term and the definition linked through Predicative Phrase (PrP) whose head is the verb *to be* connecting the term with the definition *an inflammation of the eye*. For our experiment, we focus on the identification of those NPs, within this class definitions, that introduce a Genus Term (in the example, the NP *an inflammation*), because, following to [Wilks, Slator and Guthrie, 1995], these NPs introduce good candidates to hyperornyms.

4.2 Corpus

The corpus is constituted by a set of documents of the medical domain, basically human body diseases and related topics (surgeries, treatments, and so on). These documents were collected from MedlinePlus in Spanish. MedlinePlus is a site providing information about diseases, treatments, and conditions that is easy to understand. The size of the corpus is 1.2 million of words. We chose a medical domain for reasons of availability of textual resources in digital format.

4.3 Chunking relevant fragments

We have used the Tool-Kit library of Natural Language designed by Bird, Klein and Loper [2009] in Python language, for implementing a chunker in order to extract NPs with relational adjectives and locative adverbs for the case of identifying concrete entities (reference object). In the following sections we describe the pattern considered for the chunking phase.

Noun Phrases

Given we are concerned in subordinate categories to the hypernyms (genus) extracted from analytical DCs, the structure of NPs considered is:

<NC><AQ>

Acording to Vivaldi and Rodríguez [2001], this syntactical pattern of a noun (in this case, a common noun or *NC*) and at least an adjective (that is, a *AQ*) is one of the most productive for configuring terms in Spanish.

Non-relevant adjectives

We propose a phase of reduction of *noise* in candidates. In this phase non-relevant adjectives are removed from noun phrases. Demonte [1999] proposes a set of heuristics for distinguish between relational and descriptive adjectives. As Acosta, Aguilar and Sierra [2013] point out, relational adjectives have a higher probability of being part of terms or subordinate categories. The heuristics considered in this experiment that achieved the most high precision are the following:

<RG><AQ>
<VAE><AQ>
<D.*|P.*|F.*|S.*><AQ><NC>

Where RG, AQ and VAE as tagged with FreeLing, correspond to adverbs, adjectives and the verb *estar*, respectively. Tags <D.*|P.*|F.*|S.*> correspond to determinants, pronouns, punctuation signs and prepositions. The expression <D.*|P.*|F.*|S.*> is a restriction to reduce *noise*, since elements wrongly tagged by FreeLing [Carreras *et al.*, 2004] as adjectives are extracted without this restriction.

Spacial Scenes

As mentioned on section 3.3, spatial scenes are linguistic fragments where a figure is located considering axial properties of a reference object. In Acosta and Aguilar [2015] proposed place adverbs functioning with *of* preposition in order to identify and extract the reference object that we assume to be a concrete entity. The regular expression for identifying this information is:

<RG><PDEL><DA>?<NC><AQ>*

Where <DA> is a determinant tag and <PDEL> is a tag including contraction *del* and *of* plus an article (i.e., *de la*).

4.4 Relevance of words

We applied the relative frequency ratio as in (1) between the reference and domain corpus in order to calculate relevance of words. With this goal, we take into account only nouns and adjectives because they are the most used categories in the building of terms:

$$weight(w_i) = \log_2\left(\frac{f_{w_{i,D}}}{N_{w_{i,D}}} / \frac{f_{w_{i,R}}}{N_{w_{i,R}}}\right) \quad (1)$$

Where $f_{w_{i,D}}$ and $N_{w_{i,D}}$ correspond to the absolute occurrence frequency of w_i and the size of the domain corpus, respectively. Similarly, $f_{w_{i,R}}$ and $N_{w_{i,R}}$ correspond to absolute occurrence frequency of w_i and the size of the reference corpus. Words with relative frequency in reference corpus greater than or equal domain corpus are considered as part of the stopword list. On the other hand, the relevance of words only occurring in domain corpus is calculated considering (2):

$$weight(w_i) = 1 + \log_2\left(f_{w_{i,D}}\right) \quad (2)$$

4.5 Relevance of multi-word candidates

The ranking of term candidates is done by adding up the individual ranks of words present in the candidate. Additionally, the occurrence frequency of noun phrases as a whole can be added to the sum of individual rankings. In this sense, we would expect that the higher occurrence frequency of noun phrases, the higher its syntagmatic stability. It is important to remember that before calculating ranks of multi-word candidates, these candidates have already been stripped of non-relevant information, so that the occurrence frequency of noun phrases is adjusted.

Formally, if a np (that is, a Noun Phrase) has a length of n words, $w_1\ w_2\ \ldots w_n$, where n>1, then the ranking of the candidate *np* is the sum of the frequency of *np* as a whole plus the weights of all the individual words w_i:

$$weight(w_i) = f_{np} + (\sum_{i=1}^{n} w_i) \qquad (3)$$

4.6 Reducing noise

We seek to remove non-relevant words from NPs because there is a very frequent use of descriptive adjectives with relevant information (*rare disease, serious disease*, and so on). After the chunking phase, *noise* can be reduced by removing non-relevant words such as descriptive adjectives and words (nouns or adjectives) whose relative frequency in a reference corpus is greater or equal than in the domain corpus. In the latter case we consider only nouns and adjectives of a reference corpus extracted from an on line newspaper[1] with a size of 4.5 millions of tokens.

5 Preliminary Results

Once performed our chunking process, we evaluate the sets of candidates to hyponyms, hypernyms, meronyms and holonyms obtained. In the case of hyponymy/hypernymy relations, we have obtained the following results:

Table 1. Precision, recall and F-measure for the detection of hyponyms and hypernyms

Phase	Recall	Precision	F-measure
Baseline	87%	17%	28%
Chunk Grammar	58%	62%	60%
Filter 1	57%	68%	61%
Filter 2, Freq ≥ 5	42%	80%	55%
Filter 2, Freq ≥ 6	40%	81%	54%
Filter 2, Freq ≥ 10	35%	84%	49%
Filter 2, Freq ≥ 20	27%	90%	42%

In this table we expose results of precision, recall and F-measure. Our baseline consists of extracting candidates of DCs that introduce predicative phrase whose heads were verbs as *ser*, *caracterizar*, *concebir*, *definir*, and other similars according to Sierra *et al.* [2008]. Thus, our chunk achieved a recall of 58% and a precision of 62%. Applying the first filter (filter 1) of causal relations, a precision of 68% with a reduction of recall to 57% was obtained. On the other hand, the application of frequency thresholds of occurrence of hypernyms shows best results in precision, but as thresholds are increased, recall is significantly reduced.

In the case of meronymy/holonymy relations, we have considered 5 phases of evaluation. The first phase of extraction of concrete entities had a high precision of 73%. In a next *bootstrapping* step precision decreased to 51%. Next phases (3 to 5) hold measures without major changes in precision. In relation to the candidate Part-Whole relations were only extracted with the set obtained of the phase 1. Given the low precision obtained in phase 1 (24%), candidates of remaining phases were not analyzed. In general terms, recall achieved in the extraction of concrete entities was 65%.

We think that one of the relevant results of our experiment, which was not addressed in the initial design, is the precision in recognition of terms of the domain. As it can be observed in the table 2, precision is over a 70%, which is a relevant result. On the other hand, we noted extraction of concrete entities tend to converge in a phase 5.

Table 2. Precision of candidates to concrete entities in meronymy/holonymy relations

Phase	Precision	Concrete entities	Candidates	Terms
1	73%	308	423	74%
2	51%	566	1109	74%
3	46%	744	1632	72%
4	44%	806	1849	71%
5	43%	821	1891	71%

6 Conclusions

In this paper we showed a methodology for extracting hypernyms, hyponyms, meronyms and holonyms from medical texts. For performing this extraction, we recognize first DCs following the criteria proposed by Sierra *et al.* (2008). Once obtained a set of CDs, we searched NPs involved in the lexical-semantic relations that we have mentioned. The NP pattern conformed by noun + relational adjective was one of the most productives for detecting lexical relations.

Such sequences that introduce relational adjectives specific levels of of inclusiveness between terms and Genus Terms, for example: a Genus Term as *disease* is situated in a superordinate level in contrast with *mental disease*, located in a basic level. The difference here in the relational adjective.

Genus Terms that work in a basic level can express thematic features (e.g.: *runny nose*) or classificatory features (e. g.: *nasal mucous*). Given this behavior, it is useful to recognize the type of lexical-semantic relation that is implicit in these Genus Terms.

In order to derive such relation, it is necessary to determine the status of the two elements involved in the NP, that is: the head of this NP, is a noun derived from a verb or not? An example in Spanish is: if we know that *riñón* (Eng.: *kidney*) is a concrete entity, we can infer the origin of the relational adjective *renal* (Eng.: *kidney*) in a Genus Term as *infección renal* (Eng.: *kidney infection*). Thus, we can narrow our universe of possible relations to a couple of them: locative or meronymy/holonymy relations.

In this sense, it is pertinent the development of methods for extracting automatically concrete entities. In our case, we use axial properties linked to reference objects which are involved in spatial scenes, e g.: *the*

[1] www.lajornada.com.mx

bicycle ***is in front to the house***. A NP as *mucosa nasal* (Eng.: *nasal mucous*) can represent a kind of specific mucous (hyponymy/hypernymy relation); but also the same NP can indicates that the mucous is a part of the noise (hyponymy/hypernymy relation). Of course, it would be useful useful to find a solution to this ambiguity.

Acknowledgments

This paper has been supported by the National Commission for Scientific and Technological Research (CONICYT) of Chile, Project Numbers: 3140332 and 11130565.

References

[Acosta, Aguilar and Sierra, 2013] Olga Acosta, César Aguilar and Gerardo Sierra. Using Relational Adjectives for Extracting Hyponyms from Medical Texts.In A. Lieto and M. Cruciani (eds.) *Proceedings of the First International Workshop on Artificial Intelligence and Cognition (AIC 2013)*, pages 33-44. Torino, Italy. CEUR Workshop Proceedings.

[Acosta and Aguilar, 2015] Olga Acosta and César Aguilar. Extraction of Concrete Entities and Part-Whole Relations. In B. Sharp and R. Delmonte (edts.) *Natural Language Processing and Cognitive Science. Proceedings 2014*, pages 89-100. Berlin. De Gruyter.

[Berland and Charniak, 1999] Matthew Berland and Eugene Charniak. Finding parts in very large corpora. In *Proceedings of the 37th Annual Meeting of the Association for Computational Linguistics*, pages 57-64. College Park, Maryland, USA, 1999. ACL Publications.

[Bird, Klein and Loper, 2009] Steven Bird, Ewan Klein and Edward Loper. *Natural Language Processing whit Python*. Sebastropol, Cal., 2009. O'Reilly.

[Buitelaar, Cimiano and Magnini, 2005] Paul Buitelaar, Philipp Cimiano and Bernardo Magnini. *Ontology learning from text*. Amsterdam, 2005. IOS Press.

[Carreras *et al.* 2004] Xavier Carreras, Isaac Chao, Lluís Padró and Muntsa Padró. 2004. Freeling: An open-source suite of language analyzers. In *Proceedings of the 4th International Conference on Language Resources and Evaluation*, pages: 239-242. Lisbon, Portugal. Language Resources and Evaluation Conferences.

[Ceusters and Smith, 2010] Werner Ceusters and Barry Smith. A unified framework for biomedical terminologies and ontologies. *Studies in health technology and informatics*, 160 (Pt. 2): 1050-1054, 2010.

[Croft and Cruse 2004] William Croft and David A. Cruse. *Cognitive Linguistics*. Cambridge, UK., 2004, Cambridge University Press.

[Demonte, 1999] Violeta Demonte. El adjetivo. Clases y usos. La posición del adjetivo en el sintagma nominal. In I. Bosque and V. Demonte (comps.) *Gramática descriptiva de la lengua española*, vol. 1, ch. 3, pages 129-215, Madrid, Spain. Espasa-Calpe.

[Evans and Green, 2006] Vyvyan Evans and Melanie Green. *Cognitive Linguistics: An Introduction*. Hillsdale, New Jersey, 2006. LEA.

[Evans, 2007] Vyvyan Evans. *A Glossary of Cognitive Linguistics*. Edinburgh, UK., 2007. Edinburgh University Press.

[Fellbaum, 1998] Christiane Fellbaum. *WordNet: An Electronic Lexical Database*. Cambridge, Mass., 1998. MIT Press.

[Faure and Nedellec, 1998] David Faure and Claire Nédellec. A Corpus-based Conceptual Clustering Method for Verb Frames and Ontologies. In *Proceedings of the LREC: Workshop on adapting lexical and corpus resources to sublanguages and applications*, pages 5-12. Granada, Spain. Language Resources and Evaluation Conferences.

[Girju, Badulescu and Moldovan, 2006] Roxana Girju, Adriana Badulescu and Dan Moldovan. Automatic discovery of part–whole relations. *Computational Linguistics*, 32(1): 83-135, 2006.

[Hearst, 1992] Marti Hearst. Automatic Acquisition of Hyponyms from Large Text Corpora. In *Proceedings of the Fourteenth International Conference on Computational Linguistics*, pages 539-545, Nantes, France. ACL Publications.

[Jurafsky and Martin, 2009] Daniel Jurafsky and James H. Martin. *Speech and Language Processing: An Introduction to Natural Language Processing, Speech Recognition, and Computational Linguistics*. Second Edition. New York, 2009. Prentice-Hall.

[Pantel and Pennacchiotti, 2006] Patrick Pantel and Marco Pennacchiotti. Espresso: Leveraging Generic Patterns for Automatically Harvesting Semantic Relations. In *Proceedings of the 21st International Conference on Computational Linguistics and 44th Annual Meeting of the Association for Computational Linguistics*, pages 113-120. Sydney, Australia, 2006. ACL Publications.

[Pustejovsky, 1995] James Pustejovsky. *The Generative Lexicon*, Cambridge, Mass., 1995. MIT Press.

[Rosch, 1978] Eleanor Rosch. Principles of categorization. In E. Rosh and B. Lloyd (eds.) *Cognition and Categorization*, pages 27-48. Hillsdale, New Jersey. LEA.

[Sierra *et al.*, 2008] Gerardo Sierra, Rodrigo Alarcón, César Aguilar and Carme Bach. Definitional verbal patterns for semantic relation extraction, *Terminology*, 14(1): 74-98, 2008.

[Smith and Medin, 1981] Edward Smith and Douglas Medin. Categories and Concepts. Cambridge, Mass., 1981, Harvard University Press.

[Smtih, 2004] Barry Smith. Beyond concepts: ontology as reality representation. In A.Varzi and L. Vieu (eds.) *Formal Ontologies in Information Systems*, pages 73-84. Amsterdam, 2004, IOS Press.

[Snow, Jurafsky and Ng, 2006] Rion Snow, Dan Jurafsky, and Andrew Y. Ng. Semantic Taxonomy Induction from Heterogenous Evidence. In *Proceedings of the 21st International Conference on Computational Linguistics and 44th Annual Meeting of the Association for Computational Linguistics*, pages 801–808. Sydney, Australia, 2006. ACL Publications.

[Tsujii and Ananiadou, 2005] Junichi Tsujii and Sophia Ananiadou. Thesaurus or Logical Ontology, Which One Do We Need for Text Mining? *Language Resources and Evaluation*, 39(1): 77-90, 2005.

[Vivaldi and Rodríguez, 2001] Jorge Vivaldi and Horacio Rodríguez. Improving term extraction by combining different techniques. *Terminology* 7(1): 31–48, 2001.

[Wilks, Slator and Guthrie, 1995] Yorick Wilks, Brian M. Slator and Louise M. Guthrie. *Electric Words*, Cambridge, Mass., 1995. MIT Press.

[Winston, Chaffin and Herrmann, 1987] Morton E. Winston, Roger Chaffin and Douglas Herrmann. A taxonomy of part-whole relations. *Cognitive science* 11(4): 417-444, 1987.