

Learning patterns for building resources about semantic relations in the medical domain

Mehdi Embarek and Olivier Ferret

CEA, LIST, Laboratoire d'Ingénierie de la Connaissance Multimédia Multilingue
18 route du Panorama, BP6, Fontenay-aux-Roses, F-92265, France
olivier.ferret@cea.fr

Abstract

In this article, we present a method for extracting automatically from texts semantic relations in the medical domain using linguistic patterns. These patterns refer to three levels of information about words: inflected form, lemma and part-of-speech. The method we present consists first in identifying the entities that are part of the relations to extract, that is to say diseases, exams, treatments, drugs or symptoms. Thereafter, sentences that contain couples of entities are extracted and the presence of a semantic relation is validated by applying linguistic patterns. These patterns were previously learnt automatically from a manually annotated corpus by relying on an algorithm based on the edit distance. We first report the results of an evaluation of our medical entity tagger for the five types of entities we have mentioned above and then, more globally, the results of an evaluation of our extraction method for four relations between these entities. Both evaluations were done for French.

1. Introduction

This article takes place in the medical domain, which is characterized by a rich and specific vocabulary. This peculiarity has led to the development of a large number of terminological resources in this domain as for instance the MeSH (Medical Subject Heading) or the UMLS (Unified Medical Language System). These resources were used in various tasks: document indexing, information retrieval, information extraction or question-answering.

Similarly to more general lexical networks such as WordNet (Fellbaum, 1998), these resources contain lots of paradigmatic relations like synonymy or hypernymy relations but lack of syntagmatic relations such as the ones that underlie the fact that the disease D can be cured by the treatment T or that the exam E is used to detect the disease D . Most of the methods for extracting semantic relations from texts also focus on synonymy or hypernymy relations, following (Hearst, 1992) or more recently (Caraballo, 1999).

The work we present in this article tackles the problem of extracting and validating semantic relations between entities of the medical domain, such as diseases, drugs or exams, and concentrates more particularly on syntagmatic relations. Some work was already carried out concerning the extraction of semantic relations in the medical or biomedical domain, such as (Craven, 1999), (Mukherjea and Sahay, 2006), (Rosario and Hearst, 2004), (Vintar and Buitelaar, 2003) or (Huang et al., 2006). Work focusing on information extraction was also achieved in the same domains. While its aim is supposed to be more general, it actually consists in many cases in extracting the same kind of relations, as the interactions between genes or between genes and proteins. (Nédellec, 2004) gives an overview of such work, which is often based on hand-coded rules.

Following (Pantel et al., 2004), the method we propose relies on the learning and the application of linguistic patterns that are specific to the relations to extract. The extraction of new relations is performed in two steps. The first step consists in identifying the entities from the medical domain that are part of the target relations. For instance,

in the sentence “*European Medicines Agency (EMA) has approved Revatio to treat pulmonary arterial hypertension ...*”, this first step aims at recognizing *Revatio* as a drug and *pulmonary arterial hypertension* as a disease. Then, the application of the pattern $\langle \text{Drug} \rangle * \text{to treat} \langle \text{Disease} \rangle$, that was previously learnt from a reference corpus, validates the presence of a relation between these two entities and finally leads to extract the relation $[\text{Revatio}] - (\text{cure}) - [\text{pulmonary arterial hypertension}]$ from the sentence. In the next sections, we will detail these two steps and present their application and their evaluation for French.

2. Ontology of the medical domain

The starting point of our work was the definition of an ontology that is representative of the most common entities and relations that are used by family doctors, without completeness as a goal. This ontology was built both by interviewing doctors and analyzing the questions they typically ask (Ely et al., 1999). It is made of around 20 entity types and 30 relation types.

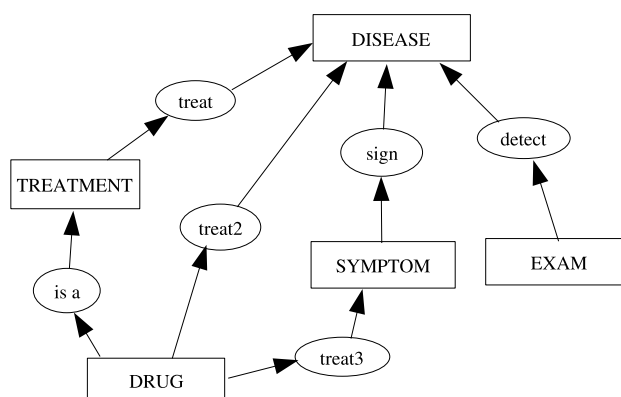


Figure 1: Subset of our ontology of the medical domain focused on the selected entities

Figure 1 shows the subset of 5 entity types, *Disease*, *Exam*, *Treatment*, *Symptom* and *Drug*, that were selected to per-

form our experiments as well as the following 4 relation types between them:

- Treat: Disease – Treatment
- Cure: Disease – Drug
- Detect: Disease – Exam
- Sign: Disease – Symptom

3. Medical entity tagging

The first step both for learning relation patterns and extracting relations consists in recognizing the medical entities of our ontology in documents. This is done in our case by applying a classical approach for recognizing general named entities, whether entities are actually named entities or complex terms: for each type of entities to identify in documents, a set of rules relying on morpho-syntactic features were defined manually from a corpus and then compiled into finite-state automata. As morpho-syntactic patterns are less significant for medical entities than for general named entities such as persons, we emphasized, for building these rules, the collection of lists of entities or parts of entities from various medical sources, such as medical Web sites¹ or dictionaries from the French Medicine Academy. Around 150 rules were developed for the recognition of the five selected types of entities.

Each rule is made of four parts: a *trigger*, the *left context* of the trigger, its *right context* and the *type* of the recognized entity. The trigger and its contexts are morpho-syntactic patterns that can refer to three levels of information about words: inflected form, part-of-speech tag and lemma. We applied the LIMA (L1c2m Multilingual Analyzer) linguistic analyzer (Besançon and de Chalendar, 2005) to documents for producing these three levels of information². The following rule

```
@DiseaseTrigger::$_L_DET?::$_L_DET ($_L_NC|$_L_NP)::
    DISEASE3
    trigger::right_context::left_context::entity_type
```

identifies *maladie de Lyme/Lyme disease*⁴ as a disease in *La maladie de Lyme est une .../The Lyme disease is a .../* while the rule

```
[@SymptomTrigger]:::[,] [$_L_NC] [$_L_DET] $_L_NC::
    SYMPTOM5
```

¹For instance, Doctissimo (<http://www.doctissimo.fr>) for drugs or Orphanet for diseases (<http://www.orphanet.net>).

²LIMA is a multilingual linguistic analyzer developed at CEA LIST that includes among others a part-of-speech tagger, a syntactic parser and a named entity recognizer

³? marks an optional element while (|) is an alternative. *\$_L_DET*, *\$_L_NC* and *\$_L_NP* are part-of-speech tags, which respectively correspond to article, noun et proper name.

⁴The translation of lexical items in examples or patterns are given as */translated lexical item/*

⁵[] marks the elements that enable the identification of an entity but are not part of it.

recognizes *fièvre/fever/* as a symptom in *...symptôme, comme la fièvre .../... symptom, as fever .../*.

The lists of entities or parts of entities we have mentioned above are present in rules through references to sets of linguistic elements having the same role, such as the elements that indicate a disease (*@DiseaseTrigger={disease, syndrome ...}*) or a symptom for instance (*@SymptomTrigger={sign, symptom ...}*).

4. Semantic relation extraction

4.1. Learning of extraction patterns from text

As mentioned in the introduction, the extraction of semantic relations is performed in our case by applying linguistic patterns. These patterns are regular expressions that can refer to three levels of information about words: inflected form, part-of-speech tag and lemmatized form. As a consequence, they are called multilevel patterns. For each type of relations, they are learnt by applying the algorithm defined by Ravichandran (Pantel et al., 2004) to a set of sentences in which this type of relations occurs. Starting from an unannotated corpus, the overall procedure for a relation type $X - (R) - Y$ is the following:

1. identification of medical named entities of types X and Y (see Section 3);
2. extraction of all sentences that contain both a X and a Y ;
3. manual selection of the sentences in which a relation of type R is actually present;
4. application of a linguistic analyzer to each selected sentence for getting the three levels of information about words. As for medical entity tagging, we used the LIMA analyzer for this task;
5. generalization of medical entities, by replacing them by their type;
6. application of the multilevel pattern induction algorithm of Ravichandran (see Algorithm 1) between each couple of selected sentences, or more specifically between the parts of the selected sentences that are delimited by X and Y ;
7. filtering of patterns.

The algorithm for inducing the most specific pattern from two sentences first computes the minimal edit distance between the two sentences by determining the minimal number of edit operations (insertion, deletion and substitution) that are necessary to transform one sentence into the other one (see the first part of Algorithm 1). The results of this first step is then used to determine an optimal alignment between the two sentences. The classical algorithm for achieving such alignment is enhanced for enabling a match of two words at one of the three available levels of information when two words are tested for a substitution (see the second part of Algorithm 1). Finally, patterns are built by completing alignments with two wildcard operators when it is necessary: (**s**) represents 0 or 1 instance of any word while (**g**) represents exactly 1 instance of any word. For

Algorithm 1 Pattern-learning algorithm (Pantel et al., 2004)

Algorithm for calculating the minimal edit distance between two sentences

Let x and y be two sentences of lengths n and m words and $D[i, j]$, the minimal edit distance between the two substrings $x[1..i]$ and $y[1..j]$.

```
D[0, 0] = 0
for i = 1 to n do D[i, 0] = D[i - 1, 0] + del(x[i])
for j = 1 to m do D[0, j] = D[0, j - 1] + ins(y[j])
for i = 1 to n do
  for j = 1 to m do
    D[i, j] = min(D[i - 1, j - 1] + subs(x[i], y[j]),
                  D[i - 1, j] + del(x[i]),
                  D[i, j - 1] + ins(y[j]))
```

print ($D[n, m]$)

where $ins()$, $del()$ and $subs()$ are the cost functions for insertion, deletion and substitution operations.

Algorithm for retrieving an optimal pattern

Let $x_1[1..i]$, $x_2[1..i]$ and $x_3[1..i]$ be respectively the level 1 (inflected form), the level 2 (lemma) and the level 3 (part-of-speech) of representation of the substring $x[1..i]$ of the sentence x . Similarly, $y_1[1..i]$, $y_2[1..i]$ and $y_3[1..i]$ are the three levels of representation of the substring $y[1..j]$ of the sentence y .

$i = n; j = m$

```
while (i ≠ 0 and j ≠ 0) do
  if (D[i, j] = D[i - 1, j] + del(x[i]))
    print (*s*)
    i = i - 1
  else if (D[i, j] = D[i, j - 1] + ins(y[j]))
    print (*s*)
    j = j - 1
  else if (x1[1..i] = y1[1..j])
    print x1[1..i]
    i = i - 1; j = j - 1
  else if (x2[1..i] = y2[1..j])
    print x2[1..i]
    i = i - 1; j = j - 1
  else if (x3[1..i] = y3[1..j])
    print x3[1..i]
    i = i - 1; j = j - 1
  else
    print (*g*)
    i = i - 1; j = j - 1
```

discarding patterns that result from the generalization of sentences that are too different from each other, only patterns whose number of wildcards was lower than 3 were kept. Moreover, patterns were sorted according to their frequency and only the 50 most frequent ones were selected.

We applied the procedure we have described above to the medical corpus of the EQueR evaluation campaign for question-answering systems in French (Ayache et al., 2006) and built patterns for the four types of relations we

selected for this study. Here are some examples of these patterns:

Disease – Exam

```
<exam> en suspicion de/in suspicion of/ <disease>
<exam> pour le/for the/ NC_GEN6 (*g*) <disease>
<disease> être/to be/ (*g*) à le/to/ <exam>
<exam> montre un/shows/ <disease>
<exam> (*g*) le diagnostic/the diagnosis/ (*g*)
<disease>
<disease> , (*s*) <exam>
```

Disease – Drug

```
<drug> est un médicament utilisé dans le traitement de
la/is a drug used in the treatment of/ <disease>
<drug> est indiqué dans le traitement de la/is indicated for
the treatment of/<disease>
<disease> (*s*) traitée par/treated by/ <drug>
<disease> chez les/for the/ NC_GEN traité par/treated by/
<drug>
<drug> ( proposé dans le traitement de/( proposed in the
treatment of/ (*s*) <disease>
<drug> dans le cas de/in case of/ <disease>
```

Disease – Symptom

```
<symptom> (*s*) , <disease>
<disease> , se manifeste par une/appears by/ <symptom>
<symptom> ( <disease>
<symptom> (*g*) être/to be/ (*s*) des symptômes d'
une/symptoms of/ <disease>
<disease> *g* avec/with/ <symptom>
<disease> VERBE_PRINC_INDICATIF7 (*s*) un/a/
<symptom>
```

Disease – Treatment

```
<treatment> dans le traitement des/in the treatment of/
<disease>
<treatment> contre le/against/ <disease>
<treatment> être/to be/ (*g*) PREP_GENERAL7 le
traitement de le/the treatment of/ (*s*) <disease>
<treatment> est recommandé pour le traitement des/is
recommended in the treatment of/ <disease>
<disease> nécessitant un/requiring/ <treatment>
<disease> , (*g*) NC_GEN (*g*) une/a/ <treatment>
```

4.2. Extraction and validation of semantic relations

For extracting new semantic relations from a corpus, *i.e.* new couples of medical entities linked by a known type of relation, we follow a two-step process. As for the learning of relation patterns, we first select candidate relations by picking sentences in which at least one couple of medical entities that are compatible with the target relation types is present. Then, we match the sentence of each candidate relation with the patterns learnt for the target relation type.

⁶NC_GEN: general common noun

⁷VERBE_PRINC_INDICATIF: main sentence verb, indicative mode; PREP_GENERAL: general preposition

If at least one pattern matches the sentence, the candidate relation is validated. Otherwise, it is discarded. More precisely, the process is the following:

1. identification of medical named entities of types *X* and *Y*;
2. extraction of all sentences that contain both a *X* and a *Y*;
3. application of a the LIMA linguistic analyzer to each selected sentence;
4. replacement of medical entities by their type;
5. for each sentence, computation of the minimal edit distance⁸ between the sentence and all the multilevel patterns for the target relation. If this distance is equal to 0 for at least one pattern, *i.e.* the relation between the two entities matches a pattern, the candidate relation is validated.

We applied this process to extract semantic relations for the four relation types of our study from a corpus of medical documents resulting from the Technolanguag ATONANT project about medical information systems. Here are some examples of relations extracted in such a way (followed by the pattern used for their validation, after \Rightarrow):

Disease – Exam

radiographie pulmonaire pour le diagnostic de *tuberculose* /pulmonary radiography for the diagnosis of tuberculosis/
 \Rightarrow <exam> (*g*) le diagnostic/the diagnosis/ (*g*) <disease>

radiographie pulmonaire pour le diagnostic de *tuberculose*/pulmonary radiography for the diagnosis of tuberculosis/
 \Rightarrow <exam> (*g*) le diagnostic/the diagnosis/ (*g*) <disease>

Disease – Drug

Insuffisance rénale chronique traitée par *Eprex* /Chronic renal insufficiency treated by Eprex/
 \Rightarrow <disease> traitée par/treated by/ <drug>

Le *vaccin* utilisé pour prévenir la *fièvre aphteuse*/The vaccine used for preventing from the foot-and-mouth disease/
 \Rightarrow <drug> utilisé pour/used for/
 VERBE_PRINC_INFINIT⁹ (*g*) <disease>

Disease – Symptom

L'*intoxication* peut provoquer des *vomissements* /The poisoning can cause vomiting/
 \Rightarrow <disease> peut/can/ VERBE_PRINC_INF DET_ART¹⁰<symptom>

Botulisme , se manifeste par une *sécheresse de la bouche*./Botulism, appears by having a dry mouth/

\Rightarrow <disease>, se manifeste/appears/ (*s*) par une/by a/ <symptom>

Disease – Treatment

chimio prophylaxie contre la *malaria* /chemioprophylaxy against malaria/
 \Rightarrow <treatment> contre le /against/ <disease>

radiothérapie dans le traitement de la *resténose*/radiotherapy for the treatment of restenose/
 \Rightarrow <treatment> dans le traitement de la/for the treatment of/ <disease>

5. Results and evaluation

We evaluated both the tagging of medical entities and the extraction of semantic relations between them following the principles that are widely applied for evaluating these two tasks. In both cases, the evaluation corpus was made of documents in French that had been downloaded from the CISMef site¹¹, a health gateway. For the tagging of medical entities, it was a subset of the medical corpus of the EQueR evaluation campaign for question-answering systems in French (1.5 MB; around 130,000 words). This evaluation corpus was manually annotated for the five medical entities of our study. Table 1 shows the results in terms of precision and recall of our medical entity tagger on this corpus. These classical measures are defined as follows in our context:

- precision is equal to the proportion of the medical named entities extracted by our system that are correct;
- recall is equal to the proportion of entities that are correctly extracted by our system among all medical entities in the corpus.

The F1-measure is the harmonic mean of precision and recall. Globally, precision and recall are at least equal to 84% on average, which is comparable to the results of the best named entity taggers for general entities such as persons or locations: the F1-measure for the best system of the CoNLL shared task on named entity identification (Tjong et al., 2003) was around 88% for English and 72% for German. Moreover, the recall of our tagger can certainly be improved by completing the resources it relies on.

For evaluating the extraction of semantic relations, we used the whole medical EQueR corpus for learning relation patterns (around 16 million words) and applied them to a subset of the corpus built for the Technolanguag ATONANT project (65 MB; around 10 million words) for extracting new relations. Contrary to what was done for medical entities, the manual annotation of the evaluation corpus for relation extraction was not done by inspecting the whole corpus but only by judging if candidate relations, *i.e.* sentences that contain a couple of entities that are compatible with a relation, were true relations. As a consequence, only the validation of candidate relations is evaluated here. Once again, we used precision and recall as evaluation measures, defined as follows:

⁸It is actually an extension of the edit distance as a match between two words can be found not only between their inflected form.

⁹VERBE_PRINC_INF: main sentence verb, infinitive mode

¹⁰DET_ART: article

¹¹<http://www.cismef.org>

| Medical entities | Precision | Recall | F-measure |
|------------------|-------------|-------------|-------------|
| Disease | 0.95 | 0.80 | 0.86 |
| Symptom | 0.84 | 0.76 | 0.79 |
| Exam | 0.94 | 0.93 | 0.93 |
| Treatment | 0.86 | 0.81 | 0.83 |
| Drug | 0.93 | 0.88 | 0.90 |
| <i>Mean</i> | <i>0.90</i> | <i>0.83</i> | <i>0.86</i> |

Table 1: Results of the tagging of medical entities for the EQueR medical corpus

- precision is equal to the proportion of extracted relations that are correct;
- recall is equal to the proportion of relations that are correctly extracted by our system among all annotated relations of the corpus.

Table 2 shows results for two kinds of patterns. Multi-level patterns are the patterns we have described above with three levels of information about words while bilevel patterns were built without taking into account the part-of-speech tags of words. Globally, the precision measure for this evaluation is quite high, especially for bilevel patterns. The lower level of the recall measure can certainly be explained both by the misses of the medical entity tagger and the fact that our relation patterns are quite specific, as their generalization is limited. Table 2 also shows quite clearly that using an abstract information such as part-of-speech tags is a good way for improving recall, even if it causes precision to decrease a little.

The comparison with similar work is somewhat difficult because types of relations, corpora and approaches are generally different. (Pustejovsky et al., 2002) used manually-designed linguistic patterns for extracting inhibitory relations from Medline with 94% for precision and 58.9% for recall. The results of our system are quite similar while our patterns were built automatically. Relation extraction is sometimes considered as a classification task applied to candidate relations. This approach was adopted by (Craven, 1999) where a Naive Bayes classifier is used to validate *subcellular-location* relations extracted from Medline with 78% as precision and 32% as recall. In the case of (Rosario and Hearst, 2004), a classifier is applied for determining the type of a candidate relation between a treatment and a disease among eight possible types. Several kinds of classifiers were evaluated and the best results – 96.9% for precision – were obtained by a neural network. Table 2 shows that we obtain results whose level is only a little bit lower than the results of (Rosario and Hearst, 2004) while we don't use resources such as a shallow parser or the MeSH database.

6. Related work and discussion

The methodology we have proposed for extracting semantic relations in the medical domain first identifies in the documents the entities that are representative of this domain, then extracts candidate relations based on the co-occurrence of these entities and finally, validates these candidate relations by applying linguistic patterns. Using such patterns for extracting semantic relations from texts is not

new. Hearst (Hearst, 1992) was among the first researchers to propose a methodology based on the acquisition and the use of lexico-syntactic patterns to extract hypernymy relations. However, the main step of the method, which consists in finding the commonalities among the environments of candidate relations for building relation patterns, was done manually. This approach was developed further by other work that also aimed at extracting semantic relations from texts and focused more particularly on building relation patterns automatically. The method we rely on and that was initially proposed in (Pantel et al., 2004) follows this trend. Moreover, the methodology of Hearst was especially successful for extracting relations in specialized domains, as it was illustrated by (Finkelstein-Landau and Morin, 1999) or (Séguéla, 1999) for instance for technical texts.

Although the method we have presented in this article is globally comparable to all these works, it differs from them by the way it applies linguistic patterns for validating candidate relations. Instead of matching candidate relations with patterns as it would be done with regular expressions, the matching is performed by computing a distance between the sentence that contains the candidate relation to validate and the validation pattern. Hence, we can implement a strict matching by requiring to have a distance equal to 0¹² or a more fuzzy matching by allowing a greater value for this distance. Moreover, this kind matching can be applied in the same way whether relations are characterized by patterns, as it is done in this work, or by examples as in a Memory-Based Learning approach. The two approaches could also be mixed.

Another important difference with work such as (Pantel et al., 2004) is that results of Section 5 were obtained without any *a posteriori* filtering of extracted relations. Despite this lack of filtering, precision is quite high while recall is not too low, which can be explained in several ways. First, this extraction is done in a specialized domain and is focused on relations between entities that are specific to this domain. Second, these relations are syntagmatic relations while relations extracted in (Pantel et al., 2004) were paradigmatic relations. Finally, the relation patterns in our case are not very general since each one comes from the generalization of only a couple of examples.

7. Conclusion and future work

In this article, we have proposed a method for extracting semantic relations between entities of the medical domain.

¹²This is the case in the experiments of Section 5.

| Relations | Multilevel patterns | | | Bilevel patterns | | |
|-------------------|---------------------|-------------|-------------|------------------|-------------|-------------|
| | Precision | Recall | F1-measure | Precision | Recall | F1-measure |
| Disease-Exam | 0.92 | 0.63 | 0.74 | 0.96 | 0.42 | 0.58 |
| Disease-Drug | 0.91 | 0.59 | 0.71 | 0.95 | 0.56 | 0.70 |
| Disease-Treatment | 0.92 | 0.69 | 0.78 | 0.99 | 0.49 | 0.65 |
| Disease-Symptom | 0.90 | 0.65 | 0.75 | 0.96 | 0.43 | 0.59 |
| <i>Mean</i> | <i>0.91</i> | <i>0.64</i> | <i>0.75</i> | <i>0.97</i> | <i>0.48</i> | <i>0.63</i> |

Table 2: Results of the extraction of semantic relations from the ATONANT corpus

This method is based on the use of multilevel linguistic patterns for validating candidate relations extracted from documents. These patterns are learnt automatically from texts in a supervised way by relying on the edit distance and an algorithm for aligning parts of sentences that takes into account several levels of information about words.

An evaluation was performed with two corpora of French documents and gave promising results. The main point to improve is recall as our evaluation has shown that the linguistic patterns that were learnt don't cover all the means by which the target relations are expressed in documents. Moreover, this evaluation was achieved only with sentences in which a couple of entities that are likely to stand for a candidate relation were detected by our medical entity tagger. Even if the results of this tagger are quite good, the need to recognize two entities for a candidate relation accentuates its errors and has an indirect impact on the evaluation of relation extraction.

In order to improve both the coverage of relation patterns and the results of our medical entity tagger, we plan to implement an iterative approach that is classically used in such a situation: we will use relation patterns not only for validating new relations but also for acquiring new medical entities by applying them to sentences in which one known entity is present together with an unknown entity. Then, these new entities will be used by our entity tagger for extending its coverage and finally, new relations will be extracted for learning new relation patterns.

Another way to improve the recall of our relation extractor would be to rely on a more advanced linguistic analysis of texts to be less dependent upon the form of relations in documents. Exploiting the results of a syntactic parser is a first step on this path, either by integrating syntactic dependency relations into patterns, as in (Snow et al., 2004), or by splitting sentences into clauses to build patterns on more homogeneous chunks of texts, as it is done in (Huang et al., 2006). Another step is to exploit semantic resources of the medical domain such as the MeSH thesaurus or the UMLS meta-thesaurus. This could be done by adding an implicit semantic level to our representation of words though the use of paradigmatic relations such as synonymy or hypernymy by an extended version of the edit distance we rely on both for learning relation patterns and exploiting them to extract new relations. Finally, we also want to extend our work by applying it to a larger part of our ontology of the medical domain. This work was voluntarily restricted to four relations but the method we have tested could be applied in the same manner to the other relations of this ontology.

8. References

- Christelle Ayache, Brigitte Grau, and Anne Vilnat. 2006. EQueR: the French Evaluation campaign of Question-Answering Systems. In *5th Conference on Language Resources and Evaluation (LREC 2006)*, pages 1157–1160, Genova, Italy.
- Romarc Besançon and Gaël de Chalendar. 2005. L'analyseur syntaxique de LIMA dans la campagne d'évaluation EASY. In *12^{ème} Conférence annuelle sur le Traitement Automatique des Langues Naturelles (TALN 2005)*, pages 21–24, Dourdan, France, June.
- Sharon A. Caraballo. 1999. Automatic acquisition of a hypernym-labeled noun hierarchy from text. In *37th Annual Meeting of the Association for Computational Linguistics (ACL'99)*, pages 120–126.
- M. Craven. 1999. Learning to extract relations from medicine. In *AAAI-99 Workshop on Machine Learning for Information Extraction*, Orlando, Florida, USA.
- J.W. Ely, J.A. Osheroff, M.H. Ebell, G.R. Bergus, B.T. Levy, M.L. Chambliss, and E.R. Evans. 1999. Analysis of questions asked by family doctors regarding patient care. *BMJ*, 319:358–361, August.
- Christiane Fellbaum, editor. 1998. *WordNet: An Electronic Lexical Database*. The MIT Press.
- M. Finkelstein-Landau and E. Morin. 1999. Extracting semantic relationships between terms: Supervised vs. unsupervised methods. In *International Workshop on Ontological Engineering on the Global Information Infrastructure*, pages 71–80.
- Marti Hearst. 1992. Automatic acquisition of hyponyms from large text corpora. In *14th International Conference on Computational Linguistics (COLING)*, Nantes, France.
- Minlie Huang, Xiaoyan Zhu, and Ming Li. 2006. A hybrid method for relation extraction from biomedical literature. *International Journal of Medical Informatics*, 75(6):443–455.
- S. Mukherjea and S. Sahay. 2006. Discovering biomedical relations utilizing the world wide web. In *Pacific Symposium on Biocomputing 11*, pages 164–175.
- Claire Nédellec. 2004. Machine Learning for Information Extraction in Genomics - State of the art and perspectives. In Spiros Sirmakessis, editor, *Text Mining and its Applications: Results of the NEMIS Launch Conference*. Springer Verlag.
- Patrick Pantel, Deepak Ravichandran, and Eduard Hovy. 2004. Towards terascale knowledge acquisition. In *20th*

- International Conference on Computational Linguistics (COLING 2004)*, pages 771–777, Geneva, Switzerland.
- J. Pustejovsky, J. Castano, and J. Zhang. 2002. Robust relational parsing over biomedical literature: Extract inhibit relations. In *PSB 2002*, pages 362–373.
- Barbara Rosario and Marti Hearst. 2004. Classifying semantic relations in bioscience texts. In *42th Annual Conference of the Association for Computational Linguistics (ACL)*.
- Rion Snow, Daniel Jurafsky, and Andrew Y Ng. 2004. Learning syntactic patterns for automatic hypernym discovery. In Lawrence K. Saul, Yair Weiss, and Léon Bottou, editors, *Advances in Neural Information Processing Systems 17 (NIPS 2004)*, pages 1297–1304. MIT Press.
- P. Séguéla. 1999. Extraction de relations sémantiques entre termes et enrichissement de modèles du domaine. In *Conférence d'Ingénierie des Connaissances (IC'99)*, pages 79–88, Palaiseau, France.
- Kim Sang Tjong, F. Erik, and Fien De Meulder. 2003. Introduction to the conll-2003 shared task: Language-independent named entity recognition. In *7th Conference on Computational Natural Language Learning (CoNLL 2003)*, pages 142–143, Edmonton, Canada.
- S. Vintar and P. Buitelaar. 2003. Semantic relations in concept-based cross-language medical information retrieval. In *ECML/PKDD Workshop on Adaptive Text Extraction and Mining (ATEM)*, Germany.