

A Hybrid System for Named Entity Metonymy Resolution

*Caroline Brun *Maud Ehrmann *Guillaume Jacquet

Xerox Research Centre Europe
6, chemin de Maupertuis
38240 Meylan France

*{Caroline.Brun, Maud.Ehrmann, Guillaume.Jacquet}@xrce.xerox.com

Abstract

This paper is an extended version of (Brun et al., 2007) describing our participation to the Metonymy Resolution (task #8) at SemEval 2007. In order to perform Named Entity metonymy resolution on location names and company names, as required for this task, we developed a hybrid system based on the use of a robust parser that extracts deep syntactic relations combined with a non supervised distributional approach, also relying on the relations extracted by the parser.

Metonymy resolution

SemEval 2007 introduces a task aiming at resolving metonymy for Named Entities, more precisely for location and organization names. As stated in the metonymy resolution task guidelines (Markert and Nissim, 2007), “metonymy is a form of figurative speech, in which one expression is used to refer to the standard referent of a related one”. The two following examples illustrate the phenomenon, on location and company names:

- (1) *He was shocked by Vietnam.*
- (2) *The BMW slowed down.*

In (1), *Vietnam*, the name of a location, refers to the war that happened here, therefore to an event. In (2), *BMW*, the name of a company, stands for the vehicle manufactured by this company.

Such phenomenon can apply in a regular way, therefore, given a semantic class (e.g. location), one can specify several regular metonymic shifts (e.g. place-for-event) that instances of the class are likely to undergo.

The goal of the task is to classify location and organization names having a metonymic or a literal reading, according to precise guidelines described in (Markert and Nissim, 2007).

Description of our System

The system we develop addresses the task of Named Entity Metonymy resolution by combining a purely symbolic approach based on robust deep parsing and lexical semantic information, with a distributional method using syntactic context similarities calculated on large corpora.

Robust and Deep Parsing Using XIP

As a fundamental component of the system we designed for named entity metonymy resolution, we use the Xerox

Incremental Parser (XIP, see (Ait et al., 2002) in order to perform robust and deep syntactic analysis. Deep syntactic analysis consists here in the construction of a set of syntactic relations¹ from an input text. These relations link lexical units of the input text and/or more complex syntactic domains that are constructed during the processing (mainly chunks, see (Abney, 1991)). These relations are labeled with deep syntactic functions. More precisely, a predicate (verbal or nominal) is linked with what we call its deep subject (SUBJ-N), its deep object (OBJ-N), and modifiers.

In addition, the parser calculates more sophisticated and complex relations using derivational morphologic properties, deep syntactic properties (subject and object of infinitives in the context of control verbs), and some limited lexical semantic coding (Levin's verb class alternations, see (Levin, 1993), and some elements of the Framenet² classification, (Ruppenhofer et al., 2006)). These deep syntactic relations correspond roughly to the agent-experiencer roles that is subsumed by the SUBJ-N relation and to the patient-theme role subsumed by the OBJ-N relation, see (Hagège and Roux, 2003) and (Brun and Hagège, 2003). Not only verbs bear these relations but also deverbal nouns with their corresponding arguments.

Here is an example of an output (chunk tree and deep syntactic relations):

“Lebanon still wanted to see the implementation of a UN resolution.”

```
TOP{SC{NP{Lebanon} FV{still wanted}} IV{to see}
NP{the implementation} PP{of NP{a UN resolution}} .}
MOD_PRE(wanted,still)
MOD_PRE(resolution,UN)
MOD_POST(implementation,resolution)
COUNTRY(Lebanon)
ORGANISATION(UN)
EXPERIENCER_PRE(wanted,Lebanon)
EXPERIENCER(see,Lebanon)
CONTENT(see,implementation)
EMBED_INFINIT(see,wanted)
OBJ-N(implement,resolution)
```

¹ Inspired from dependency grammars, see (Mel'čuk, 1998), and (Tesnière, 1959).

² <http://framenet.icsi.berkeley.edu/>

Adaptation to the Task

The parser described in previous section includes a module for “standard” Named Entity recognition, but needs to be adapted in order to handle Named Entity metonymy. Following the guidelines of the SemEval task #8, we performed a corpus study on the trial data in order to detect lexical and syntactic regularities triggering a metonymy, for both location names and organization names. For example, we examined the subject relation between organizations or locations and verbs and we then classify the verbs accordingly. From this study, we draw hypothesis like “if a location name is the subject of a verb referring to an economic action, like *import*, *provide*, *refund*, *repay*, etc., then it is a place for people. We adapted our parser by adding dedicated lexicons that encode, in the format of “semantic features”, the information collected from the corpus and developed rules modifying the interpretation of the entity, for example, the hypothesis mentioned above permits to develop the following rule:

If (LOCATION(#1) & SUBJ-N(#2[v_econ],#1))³
→ PLACE-FOR-PEOPLE(#1)

We focus our study on syntactico-semantic relations like subject, object, experiencer, content, modifiers (nominal and prepositional) and attributes. We also capitalize on the already-encoded lexical information attached to verbs by the parser, like communication verbs like *say*, *deny*, *comment*, or categories belonging to the FrameNet Experiencer subject frame, i.e. verbs like *feel*, *sense*, *see*. This information was very useful since experiencers denote persons, therefore all organizations or locations having an experiencer role can be considered as organization-for-members or place-for-people. Here is an example of output⁴, when applying the modified parser on the following sentence:

“It was the largest *Fiat* everyone had ever seen”.

ORG-FOR-PRODUCT(Fiat)
MOD_PRE(seen,ever)
SUBJ-N_PRE(was,It)
ATTRIB(It,Fiat)
EXPERIENCER_PRE(seen,everyone)
SUBJATTR(IT,Fiat)
QUALIF(Fiat,largest)
DETD(Fiat,the)

In this example, the relation QUALIF(Fiat, largest) triggers the metonymical interpretation of “Fiat”, as org-for-product.

³ Which read as “if the parser has detected a location name (#1), which is the subject of a verb (#2) bearing the feature “v_econ”, then create a LOC-FOR-PEOPLE unary predicate on #1.

⁴ Only dependencies are shown.

“The availability of *Great Britain* players has boosted the medal hopes of *Glasgow Western*”.

PLACE-FOR-PEOPLE(Great Britain)
PREP_OF(availability,players)
PREP_OF(hopes,Western)
MOD_PRE(players,Great Britain)
MOD_PRE(hopes,medal)
MOD_PRE(Western,Glasgow)
SUBJ-N_PRE(boosted,availability)
OBJ-N(boosted,hopes)
LOCATION(Glasgow)

In this example, the relation MOD_PRE⁵(players, Great Britain) triggers the metonymical interpretation of “Great Britain” as place-for-people, because “players” refers to a sport team designated by the name of the country.

In order to widen the lexical coverage of our system, we decided to use the whole BNC corpus (see <http://www.natcorp.ox.ac.uk>) to collect lexical information. We parse the corpora and then extract all previously mentioned relations involving a named entity of type location or organization along with their frequencies. For example the relation PREP_OF(invasion, COUNTRY) appears 469 times in the BNC, the relation SUBJ(COMPANY, decides) appears 420 times, etc. We then filter manually the most frequent relations in order to decide if such a context is triggering a metonymy or not, and to which class it belongs. Once the words of the context are classified, they are integrated in our semantic lexicons.

This first development step is the starting point of our methodology, which is completed by a non-supervised distributional approach described in the next section.

Hybridizing with a Distributional Approach

The distributional approach proposes to establish a distance between words depending on their syntactic distribution.

The distributional hypothesis is that words that appear in similar contexts are semantically similar (Harris, 1951).

In other words, the more two words have the same distribution, i.e. are found in the same syntactic contexts, the more they are semantically close.

We propose to apply this principle for metonymy resolution. Traditionally, the distributional approach allows to group words like USA, Britain, France, Germany because there are in the same syntactical contexts:

- (1) *Someone live in Germany.*
- (2) *Someone works in Germany.*
- (3) *Germany declares something.*
- (4) *Germany signs something.*

⁵ Modifier relation where the modifier is placed before the modified element.

The metonymy resolution task implies to distinguish the literal cases, (1) & (2), from the metonymic ones, (3) & (4). Our method establishes these distinctions using the syntactic context distribution (instead of using words distribution): we group contexts occurring with the same words: the syntactic contexts *live in* and *work in* are occurring with *Germany, France, country, city, place*, when syntactic contexts *subject-of-declare* and *subject-of-sign* are occurring with *Germany, France, someone, government, president*.

For each Named Entity annotation, the hybrid method consists in using symbolic annotation if there is (§1.2), else using distributional annotation (§1.3) as presented below.

Method:

We constructed a distributional space with the 100M-word BNC. We prepared the corpus by lemmatizing and then parsing with the same robust parser than for the symbolic approach (XIP, see section 3.1). It allows us to identify triple instances. Each triple have the form $w1.R.w2$ where $w1$ and $w2$ are lexical units and R is a syntactic relation (Lin, 1998; Kilgarriff & al. 2004).

Our approach can be distinguished from classical distributional approach by different points.

First, we use triple occurrences to build a distributional space (one triple implies two contexts and two lexical units), but we use the transpose of the classical space: each point x_i of this space is a syntactical context (with the form $R.w.$), each dimension j is a lexical units, and each value $x_i(j)$ is the frequency of corresponding triple occurrences. Second, our lexical units are words but also complex nominal groups or verbal groups. Third, contexts can be simple contexts or composed contexts (for our application, one context can be composed by two simple contexts).

We illustrate these three points on the phrase “*provide Albania with food aid*”. The XIP parser gives the following triples where for example, *food aid* is considered as a lexical unit:

OBJ-N('VERB:provide','NOUN: Albania').
 PREP_WITH('VERB: provide ','NOUN:aid').
 PREP_WITH('VERB: provide ','NP:food aid').

From these triples, we create the following lexical units and contexts (in the context *1.VERB: provide. OBJ-N*, “1” mean that the verb *provide* is the governor of the relation OBJ-N):

Words	Contexts
VERB:provide	1.VERB: provide. OBJ-N
NOUN:Albania	1.VERB: provide.PREP_WITH
NOUN:aid	2.NOUN: Albania.OBJ-N
NP:food aid	2.NOUN: aid. PREP_WITH
	2.NP: food aid. PREP_WITH
	1.VERB:provide.OBJ-N+2.N:aid. PREP_WITH
	1.VERB:provide.OBJ-N+2.NP:food aid. PREP_WITH
	1.VERB:provide.PREP_WITH +2.NO:Albania.OBJ-N

We use a heuristic to control the high productivity of these lexical units and contexts. Each lexical unit and each context should appear more than 100 times in the corpus. From the 100M-word BNC we obtained 60,849 lexical units and 140,634 contexts. Then, our distributional space have 140,634 units and 60,849 dimensions.

Using the global space to compute distances between each context appear to be too consuming and would induce artificial ambiguity (Jacquet, Venant, 2005). If any Named Entity can be used in a metonymic reading, in a given corpus each Named Entity has not the same distribution of metonymic readings. The country *Vietnam* is more frequently used as an event than *France* or *Germany*, so, knowing that a context is employed with *Vietnam* allow to reduce the metonymic ambiguity.

For this, we construct a singular sub-space depending to the context and to the lexical unit (the ambiguous Named Entity):

For a given couple context i + lexical unit j we construct a subspace as follow:

Sub_contexts = list of contexts which are occurring with the word i . If there are more than k contexts, we take only the k more frequents.

Sub_dimension = list of lexical units which are occurring with at least one of the contexts from the Sub_contexts list. If there are more than n words, we take only the n more frequents (relative frequency) with the Sub_contexts list (for this application, $k = 100$ and $n = 1,000$).

We reduce dimensions of this sub-space to 10 dimensions with a PCA (Principal Components Analysis).

In this new reduced space ($k*10$), we compute the closest context of the context j with the Euclidian distance.

At this point, we use the results of the symbolic approach described before as starter. We attribute to each context of the Sub_contexts list, the annotation, if there is, attributed by symbolic rules. In this, each kind of annotation (literal, place-for-people, place-for-event, etc) is attributed a score corresponding to the sum of the scores obtained by each context annotated with this category. The score of a context i decreases in inverse proportion to its distance from the context j : $score(context\ i) = 1/d(context\ i, context\ j)$ where $d(i,j)$ is the Euclidian distance between i and j .

We illustrate this process with the sentence *provide Albania with food aid*. The unit *Albania* is found in 384 different contexts ($|Sub_contexts| = 384$) and 54 183 lexical units are occurring with at least one of the contexts from the Sub_contexts list ($|Sub_dimension| = 54\ 183$).

After reducing dimension with PCA, we obtain the context list below ordered by closeness with the given context (1.VERB:provide.OBJ-N):

Contexts	d	symp. annot.
1.VERB:provide.OBJ-N	0.00	
1.VERB:allow.OBJ-N	0.76	place-for-people
1.VERB:include.OBJ-N	0.96	
2.ADJ:new.MOD_PRE	1.02	
1.VERB:be.SUBJ-N	1.43	
1.VERB:supply.SUBJ-N_PRE	1.47	literal
1.VERB:become.SUBJ-N_PRE	1.64	
1.VERB:come.SUBJ-N_PRE	1.69	
1.VERB:support.SUBJ-N_PRE	1.70	place-for-people
etc.		

Score for each metonymic annotation of *Albania*:

→	place-for-people	3.11
	literal	1.23
	place-for-event	0.00
	...	0.00

The score obtained by each annotation type allows annotating this occurrence of *Albania* as a *place-for-people* metonymic reading. If we can't choose only one annotation (all scores = 0 or equality between two annotations) we do not annotate.

Related work

Beyond "classical" research on metonymy (mainly focusing on common nouns), identification and resolution of metonymic readings of proper names emerged as an important challenge in computational linguistics during the last few years. Actually, this task did benefit from works of analysis and definition of what it really consists in, and from various experiments showing its difficulty but also its feasibility. Concerning the analysis of the task, K. Market and M. Nissim gave an important contribution with corpus studies for location and organization names (Market and Nissim, 2002b and 2006) that enabled them to, firstly, specify reliable frameworks for metonymy annotation and, secondly, to redefine proper names metonymy resolution as a classification task (Market and Nissim, 2002a) similar to word sense disambiguation, with the difference that the objects of disambiguation are semantic classes and not yet individual words. They experimented with word sense disambiguation methods for metonymy resolution, taking into consideration features that were not used until then and that they found relevant (Market K, 2000). They experimented, with promising results, supervised machine learning algorithms (using decision list classifier) on location (Market and Nissim, 2003) and organization names (id., 2005). These works emphasize the role of context, and the importance of grammatical features, and show its possible generalization using context similarity (via a thesaurus based on similarity between words). Other works on proper names metonymy resolution are also based on machine learning techniques: Peirsman carried out comparative experiments for location names with supervised algorithms (based on Schütze's approach) and non-supervised ones (Memory-based learning), investigating the relevance of various features (Peirsman, Y. 2006). Less complex than Market and Nissim's approach, these

methods produced state-of-the-art results. Using a different annotation framework, (Poibeau 2006) also sets out to resolve metonymic readings of French location named entities using probability levels that compute the discriminative power of various features.

Evaluation and Results

In this section, we present in detail the results we obtained with the hybrid system, on the test corpus, for organization and location names.

Type	Nb. samp	accuracy	coverage	Baseline accuracy	Baseline coverage
Loc/coarse	908	0.851	1	0.794	1
Loc/medium	908	0.848	1	0.794	1
Loc/fine	908	0.841	1	0.794	1
Org/coarse	842	0.732	1	0.618	1
Org/medium	842	0.711	1	0.618	1
Org/fine	842	0.700	1	0.618	1

Table 1: Global Results

	Nb occ.	Prec.	Recall	F-score
literal	721	0.867	0.960	0.911
place-for-people	141	0.651	0.490	0.559
Place-for-event	10	0.5	0.1	0.166
Place-for-product	1	—	0	0
Object-for-name	4	1	0.5	0.666
Object-for-representation	0	—	—	—
othermet	11	—	0	0
mixed	20	—	0	0

Table 2: Detailed Results for Locations

	Nb occ	Prec.	Recall	F-score
Literal	520	0.730	0.906	0.808
Organization-for-members	161	0.622	0.522	0.568
Organization-for-event	1	—	0	0
Organization-for-product	67	0.550	0.418	0.475
Organization-for-facility	16	0.5	0.125	0.2
Organization-for-index	3	—	0	0
Object-for-name	6	1	0.666	0.8
Othermet	8	—	0	0
Mixed	60	—	0	0

Table 3: Detailed Results for Organizations

The results obtained on the test corpora are above the baseline for both location and organization names and therefore are very encouraging for the method we developed. However, our results on the test corpora are under the ones we get on the train corpora, which indicates that there is room for improvement for our methodology.

Identified errors are of different nature:

Parsing errors: as our system makes use of a robust parser, the quality of the results depends of the quality of the syntactic analysis. For example in the sentence “Many galleries in the States, England and France declined the invitation.”, because the treatment of the coordination is not correct, France is calculated as subject of declined, a context triggering a place-for-people interpretation, which is wrong here.

Mixed cases: These phenomena, while relatively frequent in the corpora, are not properly treated by our system at the moment.

Uncovered contexts: some of the syntactico-semantic contexts triggering a metonymy are not covered by the system at the moment, neither by the symbolic component nor by the distributional analysis..

Conclusion

This paper describes a system combining a symbolic and a non-supervised distributional approach, developed for solving location and organization names metonymy. Results obtained by the system in the framework of the SemEval 2007 task#8 are reported. The competition was very stimulating and enables us to work deeply on the metonymy phenomenon, which is of great interest for NLP applications. The results obtained are quite encouraging and consequently, we plan to pursue this research work in order to improve the system on the already-covered phenomenon as well as on different names entities. We also plan to apply the same methodology on French named entities.

References

- Abney S. 1991. *Parsing by Chunks*. In Robert Berwick, Steven Abney and Carol Teny (eds.). Principle-based Parsing, Kluwer Academics Publishers.
- Aït-Mokhtar S., Chanod, J.P., Roux, C. 2002. *Robustness beyond Shallowness: Incremental Dependency Parsing*. Special issue of NLE journal.
- Brun, C., Ehrmann M., Jacquet G., 2007. *A Hybrid System for Named Entity Recognition*, Proceeding of the 4th International Workshop Semeval. ACL 2007, Prague, Czech Republic.
- Brun, C., Hagège C., 2003. *Normalization and Paraphrasing Using Symbolic Methods*, Proceeding of the Second International Workshop on Paraphrasing. ACL 2003, Sapporo, Japan.
- Hagège C., Roux C. 2003. *Entre syntaxe et sémantique: Normalisation de l'analyse syntaxique en vue de l'amélioration de l'extraction d'information*. Proceedings TALN 2003, Batz-sur-Mer, France.
- Harris Z. 1951. *Structural Linguistics*, University of Chicago Press.
- Jacquet G., Venant F. 2005. *Construction automatique de classes de sélection distributionnelle*, In Proc. TALN 2003, Dourdan.
- Kilgarrieff A., Rychly P., Smrz P., Tugwell D. 2004. *The sketch engine*. In Proc. EURALEX, pages 105-116.
- Levin, B. 1993. *English Verb Classes and Alternations – A preliminary Investigation*. The University of Chicago Press.
- Lin D. 1998. *Automatic retrieval and clustering of similar words*. In COLING-ACL, pages 768-774.
- Markert, K. 2000. *Feature Integration in Metonymy Resolution*. In: Bi-Metonymy. Bielefeld, Germany, 2000.
- Markert, K. and Nissim, M. 2002a. *Metonymy resolution as a classification task*. Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP 2002), Philadelphia, USA.
- Markert, K. and Nissim, M. 2002b. *Towards a corpus annotated for metonymies: the case of location names*. Proceedings of the third International Conference on Language Resource and Evaluation (LREC 2002), Las Palmas, Spain.
- Markert, K. and Nissim, M. 2006. *Metonymic proper names: A corpus-based account*. In Stefanowitsch A. and Gries S., editors, *Corpus-based Approaches to Metaphor and Metonymy*, Mouton de Gruyter.
- Mel'čuk I. 1988. *Dependency Syntax*. State University of New York, Albany.
- Nissim, M and Markert, K. 2003. *Syntactic features and word similarity for supervised metonymy resolution*. In Proceedings of the 41st Annual Meeting of the Association of Computational Linguistics (ACL-03), Sapporo, Japan.
- Nissim, M and Markert, K. 2005. *Learning to buy a Renault and to talk to a BMW: A supervised approach to conventional metonymy*. Proceedings of the 6th International Workshop on Computational Semantics, Tilburg.
- Nissim, M and Markert, K. 2007. *SemEval-2007 Task 08: Metonymy Resolution at SemEval-2007*. In Proceedings of SemEval-2007.
- Peirsman, Y. 2006. *What's in a name? Computational approaches to metonymical location names*. In Proceedings of the Workshop on Making Sense of Sense: Bringing Psycholinguistics and Computational Linguistics Together, Trento, Italy,
- Poibeau, T. 2006. *Dealing with Metonymic Readings of Named Entities*. Proceedings of the 28th Annual Conference of the Cognitive Science Society (COGSCI 2006), Vancouver.
- Ruppenhofer, J. Michael Ellsworth, Miriam R. L. Petruck, Christopher R Johnson and Jan Scheffczyk. 2006. *Framenet II: Extended Theory and Practice*.
- Tesnière L. 1959. *Éléments de Syntaxe Structurale*. Klincksiek Eds. (Corrected edition Paris 1969).