Nominal Metonymy Processing

Boyan Onyshkevych
US Department of Defense
Fort Meade MD 20755
baonysh@afterlife.ncsc.mil

Abstract. We argue for the necessity of resolution of metonymies for nominals (and other cases) in the context of semantics-based machine translation. By using an ontology as a search space, we are able to identify and resolve metonymic expressions with significant accuracy, both for a pre-determined inventory of metonymic types and for previously unseen cases. The entity replaced by the metonymy is made explicitly available in our meaning representation, to support translation, anaphora, and other mechanisms.

1. Introduction

Lakoff and Johnson (1980) identify metonymy as "using one entity to refer to another that is related to it." Following Gibbs (1993), we distinguish metonymy from metaphor in that metonymy uses an entity to refer to another, related, entity from the same domain, whereas metaphor necessarily relies on the replacement of an entity from one domain by an entity from another conceptual domain. As has been well-established in the literature, metonymic language use is pervasive in written and spoken language. NLP efforts addressing specific corpora, such as Hobbs and Martin (1987), Stallard (1993), and MADCOW (1992), all had to address metonymic phenomena because of its high frequency. The training and test data collected for this effort (as described below) also found high volumes of metonymy in newswires in English and Spanish. Our investigation found that the vast majority of all metonymies encountered involve the substitution of one nominal by another[1], and, given the pervasive nature of the phenomenon, we believe that semantic interpretation of nominals in context unavoidably involves metonymic resolution.

2. Why Resolve Metonymy?

We find that we need to identify and resolve metonymy during the semantic analysis phase of Machine Translation for a number of reasons, given below. (Of course, some of these arguments assume that the generation component of the MT system is able to take advantage of the additional inferences and information that is provided as a result of the resolution.)

- The most compelling argument for resolving metonymy as part of the analysis process in MT is that metonymies do not necessarily translate literally into other languages. Although often they do translate felicitously, an informal investigation into the translatability of 15 examples of metonymy easily found a number of cases where a literal translation would be bizarre, misunderstood, or just ill formed. For example, in *The newspaper fired the editor in chief*, the word for *newspaper* (*shinbun*) must be rendered as *newspaper company* (*shinbunsha*) to make the example understandable in Japanese. These results are consistent with the more thorough field work in Kamei and Wakao (1992) and Wakao and Helmreich (1993) on English, Chinese, and Japanese; they cite additional examples, such as *He read Mao* being unacceptable in Chinese. Other work by Kayser (1988) and by Horacek (1996) illustrates cases where well-formed metonymies in English are unacceptable in French or German.
- In addition to the cases where literal translation of metonymy is unacceptable, there are numerous other examples where the literal translation is understandable but not fluent.
- The replaced entity may need to be available for anaphoric and other referential mechanisms. In the

1. The non-nominal metonymic examples that we encountered, such as *spend money* used as a metonym for *go shopping*, often fall deeply in the grey area between metaphor and metonymy.

utterance *The sax has the flu tonight, so the boss docked his pay*, the pronoun refers to human (the musician) that the metonym replaced. Anaphora and definite reference function in various unique ways in different languages, so resolution is necessary for fluent translation.

- Agreement mechanisms may reference not the metonymic expression, but the replaced entity, in some examples. In the saxophone example above, the pronoun agrees with the replaced musician's gender, not the metonym's. In Japanese and other languages with counters or classifiers, expressions such as *six Volvos* require the classifier for cars, not for companies.
- Since word sense disambiguation (WSD) mechanisms typically rely on sentential context in some form, unresolved metonymies can cause inaccurate resolution.

3. Framework for Metonymy Processing

The metonymy identification and resolution mechanism described here is an integral part of the overall semantic dependency structure-building process (a process that builds the interlingual meaning representation for the input text in a Machine Translation application) in our paradigm, as it is for other applications in Hobbs and Martin (1987) or Charniak and Goldman (1988), as opposed to relegating metonymy processing to an error-recovery process, as in Fass (1986b). Because it is an integral part of the word-sense disambiguation (WSD) process, we gain efficiency and unified control, which has a high payoff because of the high prevalence of metonymy in text from real corpora. The context of this work is the MikroKosmos knowledge-based MT effort; see Onyshkevych and Nirenburg (1995) for discussion of the lexicon and other knowledge in the approach, and see Mahesh *et al.* (1997) for an overview of the WSD mechanism.

Our approach to metonymy resolution for nominals relies on a fundamental observation about metonymy, namely that it reflects (conventional) semantic contiguity, as described in Gibbs (1993) or Jakobsen and Halle (1956). The premise of our approach is that relations in our ontology[1] coincide with the relations of semantic contiguity at some level, thus the task of the metonymy resolution/WSD process is to identify the nature of contiguity in each case by identifying the best path in the ontology from the candidate meaning of a word to a constraining concept (see Mahesh *et al.* (1997) for a discussion of the richness and specificity of semantic constraints in our approach, which projected an average of 15 constraints on each open-class word in our Spanish test corpus).

By relying on the ontology to capture selectional restriction features (instead of the lexicon), and by making extensive use of inheritance in the ontology, we find that we can use a very wide range of features for constraining relations; in fact, any of the 7000 concepts in the ontology can serve as constraints, and each concept has an average of 14 constrained relations. Gibbs (1983) identifies that prior context can set up a mutually-understood local referring function: "any given instance of a referring function needs to be sanctioned by a body of beliefs encapsulated in an appropriate frame". But there are infinite such local contexts that can generate locally-sanctioned referring functions (all the "ham sandwich" types of metonymies, for example), thus an unrestricted range of notions of contiguity. While we aren't able to fully make use of context at this stage of development, the metonymy resolution/WSD process can make use of any ontological relation or predicate (event) in establishing a metonymic link. So any of the 300+ (non-inventory) relations in the ontology can all be identified as the contiguity relation and establish the metonymic link, if they provide the most plausible explanation for an apparently necessary constraint relaxation (if describing the problem from an abductive inference perspective).

---

1. Our meaning representation is defined in terms of concepts in an ontology; in addition to the traditional taxonomic (IS-A) links, we have an extensive set of other relations between concepts in the ontology, selected from over 300 possible relations. Currently the ontology consists of about 7000 concept nodes, with an average of 14 (local or inherited) relations from each concept to others in the ontology. The ontology may be examined at `http://crl.nmsu.edu/Research/Projects/mikro/htmls/ontology-htmls/onto.index.html`. References for the ontology are also available at that site.

This approach allows full use of the relations defined in the ontology. If only the strict *IS-A* relations from the ontology were used, with either vertical relaxation of constraints or a relaxation utilizing a small set of topological relations over a hierarchy (such as Fass 1986, 1988), then the wealth of metonymic expressions would be unprocessable without either allowing excessive ambiguity or not recognizing numerous uninventoried examples of metonymy. The framework outlined here allows metonymic expressions to be processed by utilizing semantic constraint checking and relaxation over the full range of metonymic relations, combined with taxonomic generalization; note, however, that not all combinations of relations or arcs in the ontology identify paths of acceptable weights, that is, the arc weight mechanism allows for identifying varying degrees of acceptability of relations that comprise potential paths between filler and constraint.

Our inventory of *metonymic arcs* reflects the types of metonymic relations which have been identified, such as *PART-OF* for the Part-for-Whole metonymy, *LOCATION-OF* for the Place-for-Event metonymy, *PRODUCTS* for the Producer-for-Product metonymy, etc. Thus for each identified metonymy, the arc(s) is found in the ontology that reflects the metonymy in defining the path from the metonym to the constraint. For example, in *he drove his V8...* the constraint on what can be driven is ENGINE-PROPELLED-VEHICLE, but the candidate filler is ENGINE (of a certain type). The part is the engine, the whole is the vehicle, and the arc from ENGINE to ENGINE-PROPELLED-VEHICLE is *PART-OF*; the potential filler is the metonym, and the constraint identifies what is being replaced. Thus in Producer-for-Product, a candidate filler (such as *Chevrolet*) has a certain relation, identified by the metonymic arc (such as *PRODUCER-OF*), to the constraint, which is what is being replaced (such as an automobile).

Thus the metonymy-processing approach described below essentially consists of two steps: a) the application of the general constraint-satisfaction process (a graph search process over the ontology), and b) identification of the concept that was replaced by the metonym in the path returned by the graph search process.

Run-time processing therefore involves finding the arc or arcs in the ontology that reflect a metonymy in the source text. Metonymic arcs would be less expensive than the rest of the unmentioned arcs, but more expensive than the weights for straightforward constraint satisfaction (i.e., *IS-A* and *INSTANCE-OF*). Yet if a straightforward constraint satisfaction path is found, the metonymic paths need not be pursued, thus not adding to the computational cost. Once a metonymic relation is found by the constraint satisfaction process, the metonym needs to be represented. The metonymic relation is represented by a slot on the metonym, which is filled by an instantiation of the concept that the metonym replaces. In other words, if X-for-Y is the metonymy, X is the metonym actually used, and Y is what it replaces, then in addition to instantiating X (from the lexical trigger), we also instantiate Y, and we connect X and Y with the metonymic arc reflecting the relation. Since every relation in the ontology has an inverse, X will have a slot *FU* filled by Y, and Y will have a slot FU^{-1} which is filled by X. A specific example of this appears below.

The general problem of acquiring the necessary static knowledge to support this approach involves identifying the list of metonymic relations, establishing relations in the ontology to reflect these metonymic relations, and assigning weights to these arcs.

For some of the metonymic relations (such as Part-for-Whole), the chaining of more than one traversals of a metonymic arc (such as the *PART-OF* arc) is acceptable; for others (such as Place-for-Event), we have a state-transition-table-based mechanism, but which is not described here.

4. Metonymy Processing: An Example

For the sentence *Lynn drives a Saab*, the semantic constraint for the appropriate slot of the appropriate sense of the verb *drive* would be *ENGINE-PROPELLED-VEHICLE. Yet the potential filler *Saab* is of type (or a subtype of) *FOR-PROFIT-MANUFACTURING-CORPORATION which is a violation of the constraint. The ontological concept *FOR-PROFIT-MANUFACTURING-CORPORATION has a slot *PRODUCER-OF*, which has an "inverse" relation called *PRODUCED-BY*. The path which is found by the ontological search process is (expressed in the [SOURCE-NODE *OUTGOING-ARC* --> DESTINATION NODE] notation):

```
ONTOLOGY PATH:
    [FOR-PROFIT-MANUFACTURING-CORP417 PRODUCER-OF --> *AUTOMOBILE]
    [*AUTOMOBILE IS-A --> *WHEELED-ENGINE-VEHICLE]
    [*WHEELED-ENGINE-VEHICLE IS-A --> *ENGINE-PROPELLED-VEHICLE]
```

If **FOR-PROFIT-MANUFACTURING-CORPORATION417** were a concept in the maned entity inventory (with knowledge about Saab Scania AB), i.e., with slot/fillers such as (*NAME* $SAAB), (*PRODUCER-OF* *AUTOMOBILE *JET-AIRCRAFT), the above path would be found. But even if that world knowledge tidbit (about Saab's products) were not available, the path that the ontological search process finds is:

```
ONTOLOGY PATH:
    [FOR-PROFIT-MANUFACTURING-CORP417 PRODUCER-OF --> *ARTIFACT]
    [*ARTIFACT SUBCLASSES --> *VEHICLE]
    [*VEHICLE SUBCLASSES --> *ENGINE-PROPELLED-VEHICLE]
```

The latter path has a lesser preference (i.e., a greater cost) than the former, because of the more expensive traversed arcs (*SUBCLASSES* is always more expensive than *IS-A*), but illustrates that the mechanism is still able to identify the metonymy in the absence of the specific product information.

Once a path is found (let's assume the latter no-named-entity-inventory case), it is inspected for the appearance of a metonymic relation arc. If such an arc is found, the inverse of that arc is available in constructing the final meaning representation of the sentence. For the above example, the most specific information that is available from the path (identifiable by following *SUBCLASSES* arcs after the metonymic arc) is utilized in making an inference about the replaced metonym and instantiating an appropriate concept **%ENGINE-PROPELLED-VEHICLE460** (the TMR is our interlingua or meaning representation language):

```
TMR:
    (DRIVE435
        (AGENT   (VALUE PERSON440)) ; abbreviated of course
        (THEME
            (SEM *ENGINE-PROPELLED-VEHICLE)
            (VALUE
                (source FOR-PROFIT-MANUFACTURING-CORP417)
                (inference inference306 ENGINE-PROPELLED-VEHICLE460))))
    (PERSON440
        (NAME $LYNN))
    (inference480
        (TYPE metonymy)
        (ENGINE-PROPELLED-VEHICLE460
            (MANUFACTURED-BY
                (VALUE FOR-PROFIT-MANUFACTURING-CORP417))))
    (FOR-PROFIT-MANUFACTURING-CORP417
        (NAME (VALUE $SAAB))
        (PRODUCER-OF inference480
            (SEM *ARTIFACT))
            (VALUE ENGINE-PROPELLED-VEHICLE460)))
```

The `inference` notation used in this example is more generally available to represent inferences made by a variety of specialized mechanisms or microtheories during the course of semantic analysis. This notation preserves the original literal interpretation, while making available the replaced entity that was inferred to exist by the metonymy processing mechanism; this inferred information (in this case, the existence of a produced vehicle) satisfies the goals of metonymy resolution mentioned above.

The real challenge to this approach is when the system has no information about the word *Saab* at all. As a system heuristic, one of the most likely possibilities for an unrecognized word in noun position (particularly if we utilize the English capitalization information) is that it is a name for some named entity (i.e., (**NAMED-ENTITY239** (*NAME* "Saab"))). In fact, we can do better by relying on Name Tagging (i.e., shallow extraction) capabilities that are available for integration into MT and other NLP applications.[1] Name Tagging technology can suggest, with high reliability (93-94%) that the string represents an organization, say **ORGANIZATION 240**, in which case the path found by the ontological search process is:

ONTOLOGY PATH:

[**ORGANIZATION239** *INSTANCE-OF* --> *ORGANIZATION]
[*ORGANIZATION *PRODUCER-OF* --> *ARTIFACT]
[*ARTIFACT *SUBCLASSES* --> VEHICLE]
[*VEHICLE *SUBCLASSES* --> *ENGINE-PROPELLED-VEHICLE]

This path, albeit expensive, is found by the search algorithm; the challenge of this approach is to adjust all of the arc weights to return these weights with fairly low cost relative to other returned paths.

5. Inventory of Metonymic Relations

Although not receiving nearly as much attention in the literature as metaphor, there have been a few attempts in the various literatures to categorize metonymy into types. None of the inventories are comprehensive enough to support the population of a working ontology for use in the analysis of real-world texts. Thus the strategy used by us to build such an inventory consisted of combining multiple sources in the literature, experiments and analysis of corpora, and carefully filtering inventories of other kinds of semantic relations (e.g., syntagmatic and paradigmatic lexical relations, meaning change, cognitive meronymic classification) for relations that do reflect metonymic use of language in English.

As mentioned above, it is not possible to build an exhaustive inventory of metonymy. So although this inventory was compiled for the purpose of seeding the metonymy processing mechanism, it is augmented with the mechanism for handling novel or unexpected (i.e., uninventoried) metonymic relations and combinations (chains) of metonymic relations.

We built an inventory of metonymy types based on various sources, spanning theoretical linguistics, lexicography, cognitive science, philosophy of language, and computational linguistics, not necessarily dealing explicitly with metonymy: Apresjan (1974), Fass (1986), Kamei and Wakao (1992), Lakoff and Johnson (1980), Mel'chuk and Zholkovsky (1984), Nunberg (1978), Stern (1965), Winston *et al.* (1987), Yamanashi (1987). Our inventory consists of about 20 major categories, with another 20 subtypes.

We encountered (in various corpora) some examples which seem to fall into multiple categories: *The White House announced that...* could be either Symbol-for-Thing-symbolized or Place-for-Occupants. There is also group of alternations that reflect a semantic relation that could be arguably treated as either metonymy, regular polysemy (i.e., handled by Lexical Rules in our format or by generative processes in Pustejovsky (1995)), or derivational processes, such as Product-for-Plant or Music-for-Dance.

We need to ensure that the metonymies in the inventory mentioned above are representable by relations in the ontology, with certain metonymies weeded out for lack of productivity (often because there is only a limited possibility of examples of the metonymy, and those are diachronically lexicalized). For each metonymic relation, we identify a relation that is used in the ontology to represent the relation between the referent and the metonym (i.e., from the thing being replaced to the thing that replaces it), along with an inverse relation (which is what actually appears in the path in a filler-to-constraint search).

A potential problem with this approach is that triggering conditions may differ from the canonical metonymy, where a selectional restriction violation is a clear indicator that some kind of relaxation is neces-

1. Numerous such Name Tagging systems, with accuracy very near human, have been evaluated in the scope of the Message Understanding Conferences (MUCs) and are described in Sundheim (1995).

sary. In particular, there might not be any selectional restriction violation for some "pragmatic" metonymies, such as *I'm going to spend money this afternoon* (which, arguably, are actually metaphors).

6. Knowledge Base for Metonymy Processing

The knowledge required for processing metonymy is not specifically differentiated from the constraint satisfaction data requirements of the overall processing mechanism. Those static knowledge resources do, however, reflect ontology arcs and weights that are used for identifying and resolving metonymy. The knowledge acquisition consisted first of identifying the arcs that needed special treatment because they are used in resolving frequently-occurring metonymies, then second by setting weights for those arcs by the automated training mechanism (using simulated annealing). The latter part of the task, however, required manually building a training data set.

The example below illustrates the training data. The example from the corpus is quoted, followed by an enumeration of the metonymy categories in effect in the example. The matrix verb is the source of constraints on the metonym in this case, so the concept is listed, along with the constraint that it places on the *AGENT* role. The path given in this example needs to be matched by the ontological graph search exactly.

```
;;; "The White House said it does not know" (USA Today)
;;; Metonymy Type: PLACENAME-FOR-OCCUPANTS
;;; Metonymy Type: ROLE-FOR-PERSON
;;; "said" = ASSERTIVE-ACT
;;; ASSERTIVE-ACT.AGENT = HUMAN (Selectional constraint)
WHITE-HOUSE (HUMAN)
      (((WHITE-HOUSE -)
        (PRESIDENT OCCUPANT)
        (ELECTED-GOVERNMENTAL-ROLE IS-A)
        (GOVERNMENTAL-ROLE IS-A)
        (SOCIAL-ROLE IS-A)
        (HUMAN IS-A)))
```

The training process for the weight assignment mechanism simply produces a weight for each of the arcs represented in a manually-produced inventory of arcs, mostly reflecting the arcs (actually, the second of each pair) identified in the inventory mentioned above. In our experiment, the arc types that receive special weights are manually specified, and the training mechanism assigns weights. It would have been possible for the training mechanism to assemble the list of arcs, as well, by examining the arcs reflected in the training data; one drawback of the latter approach would have been the inability to call out specific arcs that aren't used in the training data, in expectation of their occurrence in other corpora.

First we constructed a data set which essentially reflects an opportunistic collection of metonymies, and is in no way exhaustive or reflective of the distribution of metonymies over a corpus. Weights were produced by a simulated annealing training process; the training was able to produce a set of weights that accounted for 100% of the training data. A typical set of such weights is abbreviated below:

```
IS-A                                      0.979727
SUBCLASSES                                0.762831
HAS-MEMBER                                0.787453
PRODUCER-OF                               0.779002
...
INSTANCE-OF                               1.0
NAMED-INSTANCE-OF                         1.0
*                                         0.58028
```

The last line reflects the weight used for all arcs not explicitly inventoried.

A second training set was produced more systematically from English-language newswire, specifically the February 9-11 1997 edition of USA Today (hardcopy) and the February 11 1997 edition of the on-line edition of the Mercury News.

After the ontology was augmented as required, new weights were produced by simulated annealing. The annealing run used the same annealing schedule and Cauchy cooling rate, and began by initially "heating" the temperature (by 10 complete randomizing annealing iterations) to an energy of *0.97* (in the interval *[0.0, 1.0]*). The simulated annealing run resulted in final energy of *0.0575*, or 94.25% example accuracy (percentage of example sentences correctly analyzed, as compared to metonymic link accuracy, where examples with a chain of multiple metonymies count multiple times). Of the remaining errors (i.e., metonymic relations not found by the ontological search program), one is unsolvable by the current approach. The example, *Eddie Jones had a hot hand in today's game* has no selectional constraint violations (and is, in fact, understandable and incorrectly acceptable literally).[1] Handling this type of non-literal expression is beyond the scope of the work described here, and would require a substantially different approach.

Of the other four examples that were not solvable after training, one is actually ambiguous, and the ontological search mechanism suggested a reading not supported by context: *Fujimori told Peruvian radio that*[2]... appeared in a context which suggested that he talked to the nation via radio, vs. talking to the people in charge of the Peruvian radio service, as the ontological search program suggested. Two of the other examples, *Other dinners brought in more money* and *The dinner is adding to the questions being asked about fund-raising activities*, were incorrectly analyzed as using "dinner" to refer to the people who prepared the dinner, not the people who attended the dinner (in the former case); in is unclear how to analyze the latter of these, which is complicated by ellipsis, so there is no correct answer given in the training data, resulting in an automatic failure.[3] The last of the incorrect examples, ... *will move people from welfare rolls into jobs*, also involves some metaphorical or elliptical mechanism.[4]

7. Results

A test set was produced in exactly the same way as the training set described above, from USA Today and Mercury News articles (7 March 1997 editions). The test data in Table 1reflect the first fifty metony-

Table 1. Metonymy Test Results On English Test Data

# Correct	Errors due to missing arc	Errors due to representation gap	Errors due to bad path
47/50 = 94%	0/3	1/3	2/3

mies found in the two sources (actually, many repeat metonymies of the form *X announced... X also announced...* were omitted; the inclusion of all these (easy) metonymies would have resulted in a ratio of about 95/100 for the test set). The table shows results on this test set using weights produced by training on both the training sets described above.

The texts used for training and testing for the Spanish WSD experiments (see Mahesh *et al.* 1997) were also examined for metonymies produced as part of the semantic analysis process. The results there showed, realistically, how metonymy resolution, WSD, and building semantic dependency structure (to represent the meaning of a text) are interrelated, in that many of the WSD failures correlated with (wrong)

1. This example is from CNN, dated February 9 1997, not from Mercury News or USA Today, so doesn't really belong in this data set.
2. This example due to the on-line Mercury News service, article dated 11 February 1997.
3. Both these examples due to USA Today, 9-11 February 1997 edition (hardcopy).
4. This example due to the on-line San Jose Mercury News service, article dated 11 February 1997

metonymies being produced and vice versa. Table 2 shows the cumulative counts for different categories of

Table 2. Metonymic Inferences in 5 Spanish Texts

CORRECT METONYMIC INFERENCES	Institution for PersonResponsible	23
	ObjectUsed for User	1
	Action for Entity	15
	Generic for Specific	10
	Symbol for Symbolized	1
	Product for Producer	2
	Instrument for Action	1
INCORRECT INFERENCES	wrong preposition selected	6
	conjunction problems	2
	all other TMR errors (including bad metonymy resolution and missing microtheories)	37

metonymies produced during the course of producing TMRs for the four training and one test text. The table shows that seven kinds of metonymic expressions were found in the four texts, of which Institution-for-PersonResponsible (such as companyX announce Y, or in *El grupo Roche adquirió el laboratorio ...*) was the most common, as expected, due to the nature of the texts. The Action-for-Entity type was also well-represented, often from the use of "imports" or "exports" to refer to products (since they are represented as the *THEME* of an event in the lexical semantic specifications: *las importaciones brasileñas totalizaron ...*) The table also shows a count for various classes of errors in TMR production. A number of these errors are just reflections of missing microtheories; for example, "millions of dollars" and other numerical expressions (in Spanish) cause odd TMR constructions that cause trouble when they are linked to other elements of the SDS, resulting in type mismatches and. therefore, metonymic inferences. Another class of anomaly is due to temporal expressions, for which no microtheory has been developed, and whose absence causes funny metonymic expressions to appear in the TMR. These various missing microtheories account for about half of the errors. It is difficult to pinpoint the cause of errors, so no breakdown of the error types can be produced; for example, it is difficult to determine whether bad metonymic resolution is the cause or the effect of bad WSD on open class words or prepositions, informed by a range of other knowledge sources other than the ontological graph search. Thus many of these errors, real and apparent, would be eliminated by further development of the MikroKosmos system that formed the test environment in this case, namely by developing the following microtheories: numeric expressions, temporal expressions, reification of case roles, and prepositional semantics.

The goal of these experiments was not to attempt to solve all cases of metonymy, but to identify how far this general mechanism can lead us in addressing metonymies. In fact, the results are rather promising, in terms of coverage. A handful of examples are mentioned in above and in this subsection as difficult or impossible within the framework of the approach described here; however, they seem to account for less that five percent of metonymies occurring in real corpora. Thus we have a model which, as part of a semantic interpretation mechanisms, is able to handle a significant percentage of metonymic usage cases for nominals found in our corpora.

8. Acknowledgment

This work was done in conjunction with the MikroKosmos group at the Computing Research Lab at New Mexico State University, in particular Sergei Nirenburg, Victor Raskin, Evelyne Viegas, Steven Beale, and Kavi Mahesh.

9. References

Apresjan, Yuriy (1973). "Synonymy and Synonyms", in *Trends in Soviet Theoretical Linguistics*, F. Kiefer, ed. Dordrecht Holland: D. Reidel Publishing

Apresjan, Yuriy (1974). "Regular Polysemy" in *Linguistics* vol. 142, pp. 5-32.

Charniak, Eugene and Robert Goldman (1988). "A Logic for Semantic Interpretation", in *Proceedings of the 26th Annual Meeting of the Association for Computational Linguistics* (ACL88).

Fass, Dan (1986a). "Collative Semantics: A Description of the Meta5 Program." Computing Research Laboratory, New Mexico State University. MCCS-86-23.

Fass, Dan (1986b). "Collative Semantics: Lexical Ambiguity Resolution and Semantic Relations (with Particular Reference to Metonymy). New Mexico State University's Computing Research Lab Technical Report MCCS-86-59.

Fass, Dan (1988). "An Account of Coherence, Semantic Relations, Metonymy, and Lexical Ambiguity Resolution," in *Lexical Ambiguity Resolution*, Small, Cottrell, and Tanenhaus, eds. Palo Alto: Morgan Kaufman.

Gibbs, Raymond (1993). "Process and products in making sense of tropes," in *Metaphor and Thought*, Andrew Ortony, ed. Cambridge: Cambridge University Press.

Hobbs, Jerry and Paul Martin (1987). "Local Pragmatics". SRI International Technical Note 429.

Horacek, Helmut (1996). "On Expressing Metonymic Relations in Multiple Languages", in *Machine Translation*, vol. 11:1-3, pp. 109-158.

Jakobsen, Roman and Morris Halle (1956). *Fundamentals of Language.* 'S-Gravenhage: Mouton & Co.

Kamei, Shin-ichiro and Takahiro Wakao (1992). "Metonymy: how to treat it properly in a multilingual machine translation system", in *Proceedings of the First Singapore International Conference on Intelligent Systems.* pp. 493-497.

Kayser, Daniel (1988). "What Kind of Thing is a Concept?", *Computational Intelligence*, vol. 4:2, pp. 158-165.

Lakoff, George and Mark Johnson (1980). *Metaphors We Live By.* Chicago: University of Chicago Press.

MADCOW (1992). "Multi-Site Data Collection for a Spoken Language Corpus", in *Proceedings of the DARPA Speech and Natural Language Workshop*, Harriman NY. (MADCOW - the Multi-site ATIS Data COllection Working group).

Mahesh, Kavi, Sergei Nirenburg, Stephen Beale, Evelyne Viegas, Victor Raskin, and Boyan Onyshkevych (1997). "Word Sense Disambiguation: Why Statistics When We Have These Numbers?", in *Proceedings of the International Conference on Theoretical and Methodological Issues in Machine Translation (TMI-97)*, Santa Fe NM.

Mel'chuk, Igor and Alexander Zholkovsky (1984). *Explanatory Combinatorial Dictionary of Modern Russian: Semantico-Syntactic Studies of Russian Vocabulary.* Vienna: Wiener Slavistischer Almanach.

Nunberg, Geoffrey (1978). *The Pragmatics of Reference.* CUNY Dissertation. Reprinted by Bloomington IN: Indiana University Linguistics Club.

Onyshkevych, Boyan and Sergei Nirenburg (1995). "A Lexicon for Knowledge-Based MT", *in Machine Translation*, 10:1-2

Stallard, David (1993). "Two Kinds of Metonymy", in *Proceedings of the 31st Annual Meeting of the Association for Computational Linguistics (ACL93).* Columbus, Ohio.

Stern, Gustaf (1965). *Meaning and Change of Meaning.* Bloomington: Indiana University Press. (first edition: 1931)

Sundheim, Beth, *ed.* (1995). *The Proceedings of the Sixth Message Understanding Conference (MUC-6).* San Francisco: Morgan Kaufmann.

Wakao, Takahiro and Stephen Helmreich (1993). "Translation of Metonymy in an Interlingual MT System", in *Proceedings of the Pacific Association for Computational Linguistics (PACLING)*, Vancouver.

Winston, Morton, Roger Chaffin, and Douglas Herrmann (1987). "A Taxonomy of Part-Whole Relations," in *Cognitive Science*, vol. 11, pp. 417-444.

Yamanashi, Masa-aki (1987). "Metonymic Interpretation and Associative Processes in Natural Language." in *Language and AI (Proceedings of international symposium in Kyoto 1986)*, Makoto Nagao (ed). Elsevier/North Holland, pp. 77-86