Deverbal Nouns in Knowledge Representation*

Olga Gurevich
Linguistics Department
UC Berkeley
1203 Dwinelle Hall
Berkeley, CA 94720-2650
olya.gurevich@gmail.com

Richard Crouch, Tracy Holloway King, Valeria De Paiva
Palo Alto Research Center
3333 Coyote Hill Road
Palo Alto, CA 94304
{rcrouch, thking, paiva}@parc.com

Abstract

Deverbal nouns pose serious challenges for knowledge-representation systems. We present a method of canonicalizing deverbal noun representations, relying on a rich lexicon of verb subcategorization frames, the WordNet database, a large finite-state network for derivational morphology, and a series of heuristics for mapping deverbal arguments onto the arguments of corresponding verbs.

Introduction

Deverbal nouns, or *nominalizations*, can pose serious challenges for knowledge-representation systems. Sentences (1) and (2) describe the same event of destruction, which has the same two participants in both cases. However, the event is expressed by a verb in the first case and a noun in the second case. Most syntactic parsers will recognize the verb's arguments but not those of the noun.

(1) Alexander destroyed the city in 332 BC.

(2) Alexander's destruction of the city happened in 332 BC.

In this paper, we present a method of systematically mapping the arguments of deverbal nouns to those of the corresponding verbs, relying on a rich lexicon containing verb subcategorization frames and a series of heuristics for mapping deverbal arguments onto verbal arguments.

Our approach is a kind of canonicalization of deverbal nouns, making them look like verbs for the purposes of knowledge representation and reasoning. This type of canonicalization is essential for reasoning systems that rely on semantic, and not just syntactic, information from the input. For example, such a system should be able to answer the question in (4) based on the sentence in (3).

(3) **Input:** The acquisition by US Air of America West last week rocked the financial world.

(4) **Q:** Did US Air acquire America West?
A: Yes.

Deverbal nouns are extremely common in written and spoken English. Out of about 2000 parsed sentences of the Wall Street Journal, over half contained at least one deverbal noun (see also section on coverage below).

Despite the frequency of deverbal nouns, most lexical resources currently available do not provide systematic correspondences between deverbals and verbs (for a notable exception, see NOMLEX and a comparison with it below). The correspondence between deverbal noun arguments and verbal arguments is non-trivial and has received much attention in the linguistic literature (cf. Nunes (1993)). Much of the focus in the literature has been on making generalizations about what deverbal arguments are possible given particular classes of verbs. The resulting argument structures have been found to depend on specific lexical and aspectual properties of the base verbs (e.g. *Aktionsarten*), the thematic roles corresponding to syntactic positions (e.g. agent vs. experiencer), and several other verbal properties not easily available from the current computational lexicons.

However, from the point of view of text parsing, a more relevant problem is how to create verb-like representations from sentences with nominalizations. The flow of information here is different from the discussions in the linguistic literature. In the former, verb properties are known, and the task is to derive all and only possible combinations of deverbal arguments. In the latter, it is already known how the deverbal arguments have been assigned, and the task is to match them to verb subcategorization frames. We would like to suggest that this task is possible using several argument-mapping heuristics described below, and is less prone to ambiguity than is suggested by the linguistic literature.

The mapping between deverbal nouns and verbs happens within PARC's larger text processing system, the basics of which are described in the next section.

System Background

In the PARC system conversion of text to Knowledge Representation (KR) proceeds as follows. First the XLE parser is used to parse text against a broad-coverage, hand-written English grammar (Riezler *et. al* 2002). The parse output is fed into a semantic interpreter that produces fully scoped,

*This material is based in whole or in part on work funded by the U.S. Government, and any opinions, findings, conclusions, or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the U.S. Government.

higher-order intensional logical forms of the kind familiar to traditional formal semanticists. These logical forms are then flattened to a clausal form, similar in some respects to clausal forms obtained by skolemizing and flattening first order formulas. These clauses are then passed through the mapping component described in (Crouch 2005). This converts semantic representations induced by the linguistic structure of the sentence into an abstract knowledge representation more amenable to tractable reasoning.

The entire system places an emphasis on robustness and dealing with ambiguity. All input texts produce an output KR, though sometimes only by virtue of using less precise back-off mechanisms. Ambiguities are packed into an efficient, chart-like representation that is common to all components of the system (see (Crouch 2005) for details on how packing applies to the semantics to KR rewrite rules).

Why is a mapping from semantic representations to knowledge representations required? Surely natural language, and hence natural language semantics, is the most natural knowledge representation there is? Given an ideal analysis of language and its semantics on the one hand, and an ideal analysis of knowledge representation and implementation of reasoning on the other, one might reasonably expect the two to coincide. A number of researchers have assumed no fundamental distinction between linguistic semantics and (the current state of the art in) KR. In some cases, this is a practical consequence of working on language understanding for limited domains. Linguistic analysis can be specially tailored to meet the domain, e.g. eliminating many forms of ambiguity at the outset, and/or avoiding anything other than special case analyses of inferentially troublesome higher-order or intensional constructions (Allen *et al* 1996). In other cases it is a consequence of the assumption that natural language semantics can be made to coincide with what is currently tractable in automated reasoning: either through care and ingenuity in formulating a sophisticated first-order analysis (Hobbs 1985; Blackburn *et al* 2001), or through the assumption that cases falling outside of an essentially quantifier-free first-order logic are sufficiently rare as to be insignificant (Moldovan and Rus 2001).

Even if first-order logic were sufficient for NL semantics, there is still a clash of compositionalities between semantics and KR to be overcome. Semantic representations must respect the syntactic composition of the texts from which they are derived, to achieve a general and systematic syntax-semantics mapping. Consequently, the semantic representations assigned to sentences tend to be more complex, and different, than the representations a knowledge engineer would assign on a case-by-case basis when targeting a particular knowledge base. Semantics to KR transformations include:

1. Map words / word senses onto terms in the target ontology.

2. Make meaning postulates / lexical entailments explicit in the KR. For example (5) currently gets a representation where contexts encapsulate a possible worlds-style lexical expansion of "prevent": If A prevents B, then in the actual world (which we call context0) there is an instance of A but no instance of B, whereas in a counterfactual world (context1) just like the actual world apart from the absence of A, there is a B.

(5) The technician prevented an accident.

Our representation says that there is some sub-concept of an action performed by a technician that is instantiated in context0 but uninstantiated in context1. And it says that there is some sub-concept of accident (the type of action that was actually prevented) which is uninstantiated in context0 but instantiated in context1. As discussed in (Condoravdi *et al* 2001), this analysis accounts for downward monotonicity and non-existence entailments that are problematic under various first-order accounts.

3. Canonicalize compositionally distinct but equivalent semantic representations onto the same KR. The sentences in (6) receive different compositional semantic analyses, but are both mapped to the same KR.

(6) The technician cooled the room.
The technician lowered the temperature of the room.

The mapping of deverbals onto event representations described in this paper is essentially a type of canonicalization.

4. Eliminate ontologically ill-formed analyses. For example, in "John saw a man with Mary", selectional restrictions in the ontology may rule out the parse where the prepositional phrase modifies the verb "see". We used the basic sortal restrictions provided by VerbNet (Kipper *et al.* 2000). These have been "translated" into CYC concepts, so that the information about predicate-argument structure is of high quality when CYC concepts are used. But if the CYC concepts are not available (and this unfortunately is frequent), we back off to the coarser VerbNet sortal restrictions.

5. Reformulate intensional and higher-order aspects of the semantic representation in a form more amenable for KR (Bobrow *et al* 2005).

We now turn to how our system handles deverbal nouns.

Stage 1: Identifying Deverbal Nouns

Existing lexical resources do not contain a full list of deverbal nouns with their corresponding verbs. To approximate such a list, we extracted a list of verbs from the extensive XLE lexicon and used WordNet (Fellbaum 1998) to obtain all nouns related to those verbs.

The list obtained from WordNet consists of nouns derived from verbs using overt morphology (e.g. *statement* from *state*); words that can be either verbs or nouns (e.g. *travel*); or verbs derived from nouns (e.g. *criticize* from *critic*). WordNet itself does not differentiate between these types of relations, so another method is needed to separate out just the deverbal nouns.

In order to prune the verb-noun list, we check for derivational connections between each pair using PARC's finite-state morphological transducer (Kaplan *et al.* 2004). The transducer produces *possible* morphological formations and thus by itself would over-generate. However, constrained by the actually occurring words from WordNet, it gives us the verb-noun pairs where the noun is derived from the verb.

Because of this constraint, however, we are not able to include words that can be either verbs or nouns with no added morphology, as it is impossible to determine the direction of derivation.

We further use the output of the morphological transducer to separate morphological subclasses of deverbal nouns that have different syntactic properties. The deverbal nouns are divided into the three classes below.

1. *-er* class (e.g. *writer*)

2. *-ee* class (e.g. *employee*)

3. the rest (except gerunds; see below)

As is often the case with derived words, some deverbals are highly lexicalized and no longer retain a connection to the original verb (e.g. *dinner* is not really a nominalization of *dine*). Since it is impossible to detect such cases on a large scale, we concentrate on the most frequent nouns, using the BNC frequency list (Leech, Rayson, and Wilson 2001). The 200 most frequent nouns contain 12 deverbals, of which 8 were lexicalized so that the semantic connection to the verb is no longer salient. By contrast, of the 200 least frequent nouns, 23 are derived from verbs using overt morphology, and only one of those is lexicalized. These statistics provide some justification for relying on the list of most frequent nouns in order to detect lexicalized deverbals. We thus manually remove some frequent lexicalized nouns, resulting in a total of 3713 deverbals (1784 *-er* forms, 32 *-ee* forms, and 1897 forms from other classes).

The noun-verb pairs, annotated with the class information above, are used as assertions to link lexical entries within the Unified Lexicon (Crouch and King 2005). Some example assertions are in (7).

(7) /- deverbal_assertion(destruction, destroy, null).
/- deverbal_assertion(writer, write, er).

Gerunds (deverbal nouns derived using the suffix *-ing*, e.g. *baking*) are already recognized by the XLE parser as being related to verbs. The syntactic output marks them as verb stems with the feature [+gerund]. Therefore, there is no need to include explicit deverbal assertions in this case. Many lexicalized *-ing* nominals (like *building*) are also listed in the syntactic lexicon as nouns. When these are parsed, the syntactic processor presents both the lexicalized and the productive verbal interpretations, leaving it to KR to resolve the ambiguity.

The next step is to map the noun phrases co-occurring with the deverbal nouns to arguments of the corresponding verb forms.

Stage 2 - Mapping Deverbal Arguments

Our goal is to make deverbal nouns look like event descriptions, just like verbs, for purposes of subsequent knowledge representation and reasoning. This is accomplished by turning noun-like structures into verb-like structures after syntactic and semantic parsing, within the semantics-to-KR transfer system.

The basic mechanism is as follows.

1. Check if the noun is a deverbal mentioned in the list of deverbals, or a gerund (marked by the syntactic output).

2. If so, extract subcategorization frames for the corresponding verb (only simple intransitive and simple transitive frames were used)

3. Identify potential arguments

4. Check if these meet the verb's type restrictions on arguments

5. Create an event (semantic equivalent of a verb) for the deverbal noun

6. Insert statements linking the arguments to the new event

The linguistic literature distinguishes two types of nominals: so-called 'process' nominals vs. 'result' nominals. Process nominals imply that the event (e.g. *collection*) is taking place or has taken place, and the nominal refers to the action. Result nominals, by contrast, refer to the goal or result of the process, (e.g., the set of items that results from collecting something). Thus, result nominals are less action-like, and one would not want to turn them into verb-like representations. According to (Nunes 1993), when deverbal nouns are accompanied by overt arguments, they are more likely to be of the 'process' variety, whereas the interpretation can be ambiguous when no overt arguments are present. For the current stage of our system, we have decided to default to the process interpretation in both situations. However, see section on further work for other suggestions.

Some of the types of argument mappings are described below.

Transitive Verbs

Deverbal nouns derived from transitive verbs can have both the subject and object overtly expressed; in some cases, however, one or both arguments may be implicit. Some examples of transitive argument mappings are in Table 1.

Subject	Object	*Example*
possessive	of-phrase	*Ed's destruction of the city*
possessive	for-phrase	*Ed's advocacy for freedom*
by-phrase	of/for	*The destruction of the city by Ed*
—	of-phrase	*The destruction of the city*
—	preposed	*City destruction*
—	possessive	*The city's destruction*
possessive	—	*Ed's assessment*

Table 1: Argument mappings for transitive verbs

As can be seen, the last pattern is ambiguous: when only one argument is expressed by a possessive phrase, it could be the subject or the object. In addition, the possessive phrase could refer to a temporal or spatial description that is not an argument at all, e.g. *Yesterday's destruction (of the city)*, or the noun may not be used in its deverbal meaning, so that the possessive phrase refers to a true possessor.

To account for such ambiguity, an optional conversion to an event is created, splitting the choice space in the semantics-to-KR system. In one variant, the noun is left as a noun, with the possessive argument treated accordingly; in the other variant, the deverbal noun is converted into a verb-like semantic structure.

Ed's box.
possesses(box##2, Ed##0),
subconcept(Ed##0, Person),
subconcept(box##2, BoxTheContainer).

Ed destroyed the city.
doneBy(destroy_ev##1, Ed##0),
inputsDestroyed(destroy_ev##1, city##3),
subconcept(destroy_ev##1, DestructionEvent),
subconcept(Ed##0, Person),
subconcept(city##3, City).

Ed's destruction of the city.
doneBy(destruction##2, Ed##0),
inputsDestroyed(destruction##2, city##6),
subconcept(destruction##2, DestructionEvent),
subconcept(Ed##0, Person)),
subconcept(city##6, City).

Figure 1: KR for noun, verb, and deverbal.

The verb subcategorization frames in VerbNet and CYC provide some information on semantic type restrictions relevant for the different arguments. Where possible, this information is used to resolve the ambiguities: the mapping of deverbal arguments happens after the initial semantic parsing, so their semantic types are already known. For example, the verb *destroy* prefers an animate subject and allows both animate and inanimate objects. In the phrase *the city's destruction*, the possessive phrase *city* cannot fill the subject role and must therefore be the object. By contrast, in the phrase *Ed's assessment*, Ed is much more likely to be the subject of the verb *assess*.

Cases where only one argument is expressed are treated similarly to cases of verbs with missing arguments: an implicit argument slot is created.

Figure 1 demonstrates the (simplified) flattened KR representation for a noun, a verb, and a deverbal noun.

Intransitive Verbs

If a deverbal noun is derived from an intransitive verb, its argument (the subject) can be expressed by an *of*-phrase or a possessive phrase, as in Table 2.

Subject	*Example*
possessive	*Ed's death*
of-phrase	*The death of Ed*

Table 2: Argument mappings for intransitive verbs

In either case, the deverbal noun is converted to a verb-like structure, and the possessive or of-phrase is linked to the subject role.

Many verbs in our verb lexicon have both transitive and intransitive subcategorization frames. Since the transfer system operates in a fixed order, consuming resources as it goes along, it is possible to use rule ordering to avoid spurious ambiguity problems. We first check for cases where both arguments are explicitly present and link them to the corresponding transitive verb frames. In cases where only one argument is explicitly present in the input, preference is given to intransitive verb frames. The transitive verb frames with an implicit argument are used as a last resort.

Some deverbal nouns appear in a discourse without any overt arguments. In such cases, it is very difficult to distinguish between true deverbals and nouns that are used with a purely nominal function, without any reference to an event. Currently, our system treats such nouns as deverbals and converts them into event-like structures.

-er Deverbals

The deverbal mappings created in the lexicon identify separately nouns derived with the suffix *-er*. Unlike the deverbals described in the previous sections, *-er* nominals can refer both to the event and to the subject of that event. For example, *speaker* implies the existence of the event of speaking and identifies its agent.

Something similar to the process vs. result distinction applies here as well. Process nominals imply that the event is taking place or has taken place, and the nominal refers to a participant in that event. For example, *Ed is an admirer or art* implies that the admiring has taken place, with Ed as the subject. By contrast, result nominals simply suggest that the designate of the noun is inclined to perform a certain action, but may not have ever performed it in practice. E.g., *Ed is a teacher* may mean that Ed has been trained to teach, but is not necessarily teaching now or has ever taught before. For purposes of event identification, we are only interested in the process nominals. (Nunes 1993) suggests that *-er* nominals with overtly expressed object arguments can only receive the process interpretation, whereas nominals without overt arguments could be interpreted as both. The subject can be expressed as in Table 3.

Subject	*Example*
of-phrase	*admirer of art*
preposed noun	*story teller*

Table 3: Argument mappings for transitive *-er* nominals

For *-er* nominals derived from intransitive verbs (e.g. *abdicator*), the process vs. result distinction appears to be less salient. Thus, we treated all intransitive *-er* nouns as deverbal and converted them into event-like structures.

-ee Deverbals

Deverbals derived with the suffix *-ee* consistently refer to the *object* of the action. For transitive verbs, the subject of the verb maps onto deverbal arguments as in Table 4.

Subject	*Example*
of-phrase	*employee of the company*
possessive	*the company's employee*
preposed noun	*company employee*

Table 4: Argument mappings for transitive *-ee* nominals

Coverage and Comparison to NOMLEX(-PLUS)

To find out the rate of occurrence of deverbal nouns, we parsed 2002 sentences from the Wall Street Journal corpus, with the average length of almost 20 words; of those sentences, 1087 were found to have at least one deverbal noun from our list. The total contained 867 gerunds and 1220 derivational nouns (like *destruction*). A hand-checked sample of 100 sentences revealed that the system detected most of the deverbal nouns and correctly identified their arguments. Since there are no available large hand-collected test corpora where deverbals are marked, a more wide-scale evaluation of accuracy was not possible at this time.

Another metric of evaluation is by comparison with NOMLEX (Macleod *et al.* 1998), a broad-coverage, large-scale approach to nominalizations. NOMLEX is a hand-coded database of 1000 verb nominalizations, including argument structure information. NOMLEX has been extended into NOMLEX-PLUS (Meyers *et al.* 2004), which includes about 5000 deverbal nouns, as well as de-adjectival and de-adverbial nouns.

The major differences between our system and NOMLEX and its cousins include the conceptual treatment of deverbal nouns and the amount of hand-coding needed to obtain the argument-structure correspondences.

Conceptually, NOMLEX treats deverbal nouns essentially as nouns. Nominal argument structures are considered to be related, but different from, verbal argument structures. A knowledge representation system using NOMLEX would therefore need an extra step in order to relate events denoted by deverbal nouns to those denoted by verbs. In our system, by contrast, there is no separate level of nominal argument structures; rather, deverbal nouns and their arguments are immediately converted into verb-like event structures.

NOMLEX-PLUS was created by comparing nouns and verbs that begin with the same 'prefixes', e.g. *destr*-oy and *destr*-uction. In essence, this method creates an ad-hoc derivational morphology. In our case, we used an existing FST derivational morphology, which is likely to be more accurate than the ad-hoc method. In addition, using this existing resource means less hand-coding to obtain the same results.

NOMLEX-PLUS introduces semantic divisions of the deverbal nouns, e.g. partitive vs. attributive nouns. Instead of these divisions, our system relies on derivational morphology to make generalizations about argument structures (e.g. *-er* nouns vs. *-ee* nouns). In general, both these types of distinctions are extremely useful, however, and we would like to merge the two resources to see if the two types of subdivisions can improve the accuracy of the argument mappings.

A detailed comparison of the words covered by our system and by NOMLEX-PLUS is in Table 5. About half of the NOMLEX-PLUS nominalizations are also contained in the PARC system. The instances of the same nouns being mapped to different verbs were mostly due to different spellings.

As can be seen, most of the differences in coverage are due to the different methods by which the lists of deverbal nouns were obtained. The PARC system handles many more gerunds because those are derived automatically from any verb stem, whereas the gerunds in NOMLEX-PLUS were hand-picked. Similarly, the NOMLEX-PLUS list contains deverbals that are morphologically identical to their corresponding verbs, as well as deverbals that are suppletive with respect to the verbs (such as *accolade*∼*award*), which our automatic method could not derive. Our system appears to have done substantially better than NOMLEX-PLUS in detecting *-er* nouns, assuming that our frequency-based method of excluding lexicalized nouns worked reasonably well.

Deverbal Nouns	PARC	NOMLEX+
Total nouns	13464	5034
gerunds	9751	463
non-gerunds	3175	5471
-er nouns	1784	918
-ee nouns	32	39
other	1897	3614
Overlap	2605	
Same noun, different verb	34	
In NOMLEX+ only		
Total nouns	2395	
No overt morphology	784	
Suppletion	687	
Other	924	
In PARC set only		
Total nouns	1573	
-er nouns	960	
-ee nouns	14	
Other	599	
Verb Subcategorization Frames (non-gerunds)	PARC	NOMLEX+
Unique stems	3035	2864
Total intransitive	1456	841
Total transitive	2587	2415
Intransitive only	383	195
Transitive only	1514	1769
Both	1073	646
Neither	66	254

Table 5: PARC vs. NOMLEX+ coverage.

We also compared the numbers of simple intransitive and simple transitive verb subcategorization frames included in both systems. The PARC system has a somewhat larger number of verb stems and thus, larger numbers of transitive and intransitive frames. Interestingly, the NOMLEX system has many more cases of multiple deverbal nouns derived from the same verb stem (half as many unique stems as deverbal nouns), whereas PARC deverbals map onto verb stems almost one-to-one. We do not currently have an explanation for this difference.

Further Work

Our immediate goal is to conduct further testing and coverage evaluation of our system. In particular, we would like to replace our deverbal rules with those derived from NOMLEX-PLUS to see if there are significant gains or losses in performance.

In our initial implementation, we focused on several of the more frequent classes of deverbal nouns. There are additional classes which we plan to implement in the near future. The first are nouns which take *that* clause complements. Not all such nouns are deverbal, e.g., *the idea that the world is flat*. However, for the class that are deverbal, e.g., *the statement that the world is flat*, the *that* complement corresponds to the *that* complement of the verb. As such, a very simple rule can be written mapping the deverbal noun to the verb; this rule simply relies on the assertions connecting nouns to verbs, e.g., *statement* to *state*. This correlation also provides a way to extend the syntactic coverage of the grammar. The XLE parser currently knows which nouns can take *that* complements. This list was largely extracted automatically from corpora. However, it can be extended by looking at noun-verb pairs where the verb takes a *that* complement. For such nouns, it can be hypothesized that they too can take a *that* complement, thereby allowing these forms to be parsed correctly and providing more accurate input to the semantics to KR rules.

The second class involves deverbal nouns which occur in combination with so-called ‘light verbs’ such as *have, take, get, make, find*. Light verbs do not contribute any lexical meaning to the utterance, but rather provide a semantically empty action. Their subject or object is shared with their deverbal argument, which in essence denotes the event (8).

(8) US Air *made an offer* to acquire America West.

The problem with light verbs is that their behavior is very lexically specific, and the mapping of deverbal arguments to light verb arguments is not always straightforward. Future additions to our system will classify light verbs according to which argument is shared with the deverbal, and in what syntactic contexts.

The third class consists of further divisions based on derivational morphology. Initial observations suggest that the suffix used to derive a deverbal noun from a verb may have semantic and syntactic correlates. For example, *-ment* nouns like *statement* seem to be more likely to take a *that* complement. Further study is needed to identify the accuracy and usefulness of such distinctions.

In the above discussion of ‘process’ vs. ‘result’ nominals, we mentioned that the process interpretation is more likely when a deverbal noun has overt arguments. However, there is no clear guidance in the linguistic literature for making this distinction for nouns without any overt arguments. Further research will focus on isolating semantic classes or syntactic contexts in which this distinction is easier to make.

Our approach of mapping deverbal nouns and their arguments onto their verbal counterparts and then having the usual verb semantics-to-KR rules apply to them brings up an interesting issue. Here we have described a system which maps from the string through syntax and semantics into KR. However, the same XLE technology used in parsing into KR can be used to generate strings from KR. When thinking about generation, the issue as to whether to realize a KR predicate and its arguments as a verb form or a deverbal nominal form arises. In some cases the decision is clear: each sentence must be headed by a verb; certain verbs only take nominal arguments. However, in other cases a choice can be made. For example, *I know the Romans destroyed the city.* vs. *I know of the Romans’ destruction of the city* or *His statement that the world is flat bothered me.* vs. *That he stated that the world is flat bothered me.* To produce truly natural-sounding English from KR, further study will be needed to determine when a verbal form should be used and when a nominal one is appropriate.

References

Allen, J., Miller, Ringger & Sikorski. 1996. A robust system for natural spoken dialogue. *Proc. 34th ACL.*

Blackburn, P., Bos, Kohlhase & Nivelle. 2001. Inference in computational semantics. In Bunt et al (eds) *Computing Meaning, vol 2*, Kluwer.

Bobrow, D., Condoravdi, Crouch, Kaplan, Karttunen, King, de Paiva, & Zaenen. 2005. A Basic Logic for Textual Inference. In *Proceedings of the AAAI Workshop on Inference for Textual Question Answering*.

Condoravdi, C., Crouch, van den Berg, Stolle, Paiva, Everett & Bobrow. 2001. Preventing existence. *Proc. Conference on Formal Ontologies in Information Systems*.

Crouch, D., and T.H. King. 2005. Unifying Lexical Resources. *Proceedings of the Interdisciplinary Workshop on the Identification and Representation of Verb Features and Verb Classes*, Saarbruecken, Germany.

Crouch, D. 2005. Packed rewriting for mapping semantics to KR. *Proceedings of the International Workshop on Computational Semantics*, Tilburg.

Fellbaum, C. (ed). 1998. *WordNet: An Electronic Lexical Database* MIT Press.

Hobbs, J. 1985. Ontological promiscuity. In *Proc 23rd ACL.*

Leech, G., P. Rayson, and A. Wilson. 2001. *Word Frequencies in Written and Spoken English: based on the British National Corpus.* London: Longman.

Kaplan, R., Maxwell, King, Crouch. 2004. Integrating Finite-state Technology with Deep LFG Grammars. In *Proceedings of the ESSLLI’04 Workshop on Combining Shallow and Deep Processing for NLP*.

Kipper, K., Dang, Palmer. 2000. Class-Based Construction of a Verb Lexicon. In *Proceedings of the AAAI-2000 Seventeenth National Conference on Artificial Intelligence*.

Macleod, D., Grishman, Meyers, Barrett, Reeves. 1998. NOMLEX: A Lexicon of Nominalizations. *EURALEX’98.*

Meyers, A., Reeves, Macleod, Szekely, Zielinska, Young. 2004. The Cross-Breeding of Dictionaries. in Proceedings of LREC-2004

Moldovan, D and Rus, V. 2001. Logic form transformation of WordNet and its applicability to question answering. *ACL’01.*

Nunes, M. 1993. Argument Linking in English Derived Nominals. In Van Valin (ed) *Advances in Role and Reference Grammar*, John Benjamins, 375-432.

Riezler, S., King, Kaplan, Maxwell, Crouch & Johnson. 2002. Parsing the Wall Street Journal using a Lexical-Functional Grammar and discriminative estimation techniques. *ACL’02, Philadelphia, PA.*