# Large Scale Language Independent Generation Using Thematic Hierarchies

**Nizar Habash and Bonnie Dorr**
Institute for Advanced Computer Studies
University of Maryland College Park
{habash,bonnie}@umiacs.umd.edu

**Abstract**

This paper describes a large-scale language-independent evaluation of the use of Thematic Hierarchies in natural language generation. We translate from a corpus of sentences reflecting the full variety of behavior of Levin-based verb classes. The corpus is used as input to a generation system that utilizes the same thematic hierarchy for realizing relative argument surface positions in two languages: English and Spanish. The output was manually evaluated by English and Spanish speakers. The contributions of this work include: (1) an improved thematic hierarchy over an earlier implementation; (2) a large-scale evaluation of the use of thematic hierarchies in two languages; (3) an implementation of a language independent module for natural language generation; and (4) the creation of a single tool for incremental development of multilingual lexicons.

**Keywords**

Natural Language Generation, Lexical Conceptual Structure, Thematic Hierarchy

## 1 Motivation

In (Dorr et al., 1998), an implementation of thematic hierarchies for efficient natural language generation was presented. The use of the thematic hierarchy was evaluated using a small hand-constructed corpus of 100 English sentences reflecting a variety of English verb classes and alternations. The hierarchy was implemented using cascading rules within the grammar formalism provided as part of the natural language realization engine Nitrogen (Langkilde and Knight, 1998a; Langkilde and Knight, 1998b). Some of the shortcomings of this earlier work are: (1) inadequate evaluation due to the use of a small test corpus; (2) limitation of the approach to one language only (English); (3) lack of a principled design in the implementation.

This paper presents more systematic implementation of thematic hierarchies and a large-scale evaluation of their use for generation in English and Spanish. This evaluation was helpful in incremental development of both the thematic hierarchy and the English and Spanish lexicons.

## 2 Research Context

The work presented here is part of the generation component (Traum and Habash, 2000) of the interlingual Machine Translation effort at the University of Maryland College Park. The generation component has also been used in Cross-Language Information Retrieval research (Levow et al., 2000). The interlingual representation used is Lexical Conceptual Structure (LCS), a compositional abstraction with language-independent properties that transcend structural idiosyncrasies (Jackendoff, 1983; Jackendoff, 1990; Jackendoff, 1996). This representation has been used as the interlingua of several projects such as UNITRAN (Dorr et al., 1993) and MILT (Dorr, 1997).

## 3 Overview of Generation in LCS-based Machine Translation

One of the major challenges in natural language processing is the ability to make use of existing resources. Large differences in syntax, semantics, and ontologies of such resources create significant barriers to their usage in large-scale applications. A case in point is the wide range of "interlingual representations" used in machine translation and cross-language processing. Such representations are becoming increasingly prevalent, yet views vary widely as to what these should be composed of, varying from purely conceptual knowledge-representations, having little to do with the structure of language, to very syntactic representations, maintaining most of the idiosyncrasies of the source languages. In our generation system we make use of resources associated with two different (kinds of) interlingua structures: *Lexical Conceptual Structure* (LCS), and the *Abstract Meaning Representations* (AMR) used at USC/ISI (Langkilde and Knight, 1998a). The two representations serve different but complementary roles in the translation process. The deeper lexical-semantic expressiveness of LCS is essential for language independent Lexical Selection that transcends translation divergences. The shallower yet mixed semantic-syntactic nature of AMRs makes it easier to use for target language realization.

The use of two representations in generation mirrors the use of two representations on the analysis side of the MT system, in which a parsing output is passed to a semantic-composition module; the target-language AMR is analogous to the source-language parse tree. (See Figure 1.) The Composition module takes the source-language parse tree and creates a deeper semantic representation (the LCS) using a source-language lexicon. In generation, the Decomposition module performs a reverse step that uses a target-language lexicon to create the parse-like AMR. This step is referred to as Lex-
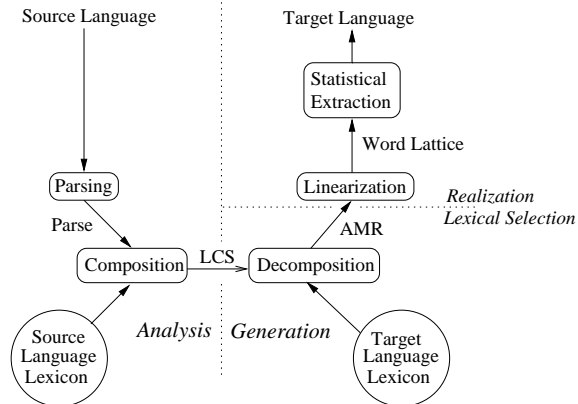
Figure 1: LCS-based Machine Translation

ical Selection. It is followed with the Realization step in which the Linearization module flattens an AMR into a sequence of words. Because of the ambiguity inherent in all of the involved modules from the parser to the lexicons, multiple sequences are created. We use the Statistical Extraction module of the generation system Nitrogen (Langkilde and Knight, 1998a; Langkilde and Knight, 1998b) to select among alternative outputs when generating English.

## 3.1 LCS Lexicons

The LCS lexicons used in both analysis and generation relate a lexeme to a Lexical Conceptual Structure representation. A single verb might have several entries corresponding to different senses of the that verb. Figure 2 compares three out of the nine root LCS (RLCS) entries for the verb 'run' in the English LCS Lexicon. These entries are classified by their Levin verb class which is used as a template to generate the RLCSes for every verb in the class. The star-marked nodes in those entries signify the location an argument can be attached. A composed LCS (CLCS) is made up of a RLCS that has its star-marked nodes filled with other CLCSes. The number at the end of the nodes mark the thematic role associated with the specific node. For example, 1 is agent, 2 is theme, 3 is a source particle (i.e. an oblique) and 4 is source (an argument). For a listing of the most common thematic roles and their corresponding codes see Figure 3. The last LCS entry for run in Figure 2 can be read as *a theme thing goes locationally from a source location to a goal location in a running manner.*

The current English verb lexicon contains over 11,000 RLCS entries. These entries correspond to different senses of over than 4,000 verbs. The Spanish verb lexicon contains over 24,000 entries corresponding to 3,300 verbs[1] The LCS lexicon also contains other information of importance to realization such as requirements for optionality (:OPTIONAL and :OBLIGATORY) and internal/external posi-

---

[1] For a detailed discussion of the acquisition of LCS-based lexicons, see (Dorr and Olsen, 1996; Dorr, 1997).

Figure 2: RLCS entries for 'run'

```
26.3 Verbs of Preparing
(cause (* thing 1)
   (go ident (* thing 2)
     (toward ident (thing 2)
       (at ident (thing 2) (run+ed 9)))))
   ((* for 17) poss (*head*) (* thing 18)))
```
Example: *John ran the store for Mary.*
Other verbs: bake boil clean cook fix fry grill iron prepare ...

```
47.5.1.b Swarm Verbs (Locational)
(act loc (* thing 2)
   ((* [at] 10) loc (thing 2) (thing 11))
   (run+ingly 26))
```
Example: *The dogs run in the forest.*
Other verbs: bustle crawl creep run swarm swim teem ...

```
51.3.2.a.i Run Verbs - (Locational,Theme only)
(go loc (* thing 2)
   ((* from 3) loc (thing 2)
             ([at] loc (thing 2) (thing 4)))
   ((* to 5) loc (thing 2)
             ([at] loc (thing 2) (thing 6)))
   (run+ingly 26))
```
Example: *The horse ran from the barn to the field.*
Other verbs: climb crawl fly jog jump leap race run swim walk ...

Figure 3: Partial List of Thematic Roles

| # | Thematic Role | Definition |
|---|---|---|
| 1 | AG | agent |
| 2 | TH, EXP | theme or experiencer |
| 3 | SRC() | source particle |
| 4 | SRC | source |
| 5 | GOAL() | goal particle |
| 6 | GOAL | goal |
| 7 | PERC() | perceived item particle |
| 8 | PERC | perceived item |
| 9 | PRED | identificational predicate |
| 10 | LOC() | locational particle |
| 11 | LOC | location |
| 12 | POSS | possessional predicate |
| 13 | TIME() | temporal particle |
| 14 | TIME | time |
| 15 | MOD-POSS() | possessional particle |
| 16 | MOD-POSS | possessed item modifier |
| 17 | BEN() | beneficiary particle |
| 18 | BEN | benefactive modifier |
| 19 | INSTR() | instrumental particle |
| 20 | INSTR | instrument modifier |
| 21 | PURP() | purpose particle |
| 22 | PURP | purpose modifier or reason |
| 25 | MANNER() | manner |
| 26 | | reserved for conflated manner |
| 27 | PROP | event or state |

tioning (:INT and :EXT). Optionality markers are necessary to determine which arguments must be available in the CLCS for proper generation using an RLCS. For example, in class 51.3.2.a.i in Figure 2, the theme is the only obligatory argument. Internal/external positioning markers will be discussed later in the section on Thematic Hierarchy.

## 3.2 Lexical Selection

The lexical selection process attempts to decompose a CLCS into RLCSes corresponding to lexemes

*John broke into the room*



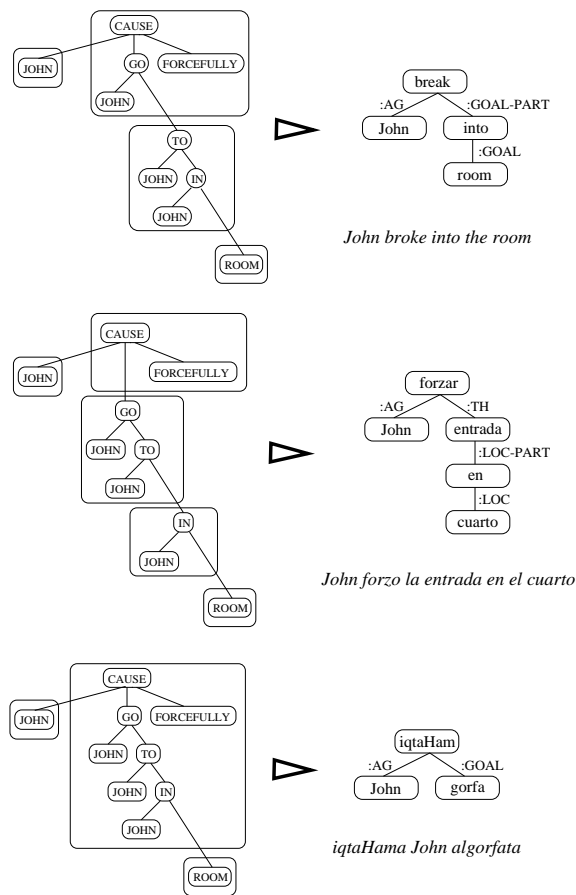*John forzo la entrada en el cuarto*



*iqtaHama John algorfata*

Figure 4: Different CLCS Decompositions

in the target language. Decomposition is basically a complex algorithm for graph matching/covering with restrictions. Its output is the shallower Abstract Meaning Representation (AMR) discussed earlier. Different lexicons for different languages provide different RLCSes and RLCS restrictions that guide lexical selection. Figure 4 compares three different possible decompositions for a CLCS into English, Spanish and Arabic. The CLCS can be read as *John causes himself to go into a room in a forceful manner*. The AMR relation (:AG, :TH, etc.) marking the connections on the left-hand side in Figure 4 are created from the thematic role information in the RLCSes.

### 3.3 Realization

Syntactic realization is the step that converts the unordered dependency tree structure of an AMR into a surface sentence. There are two operations involved in realization: recasting and linearization. Recasting converts an AMR node into another AMR node with added information, deleted information or just modified information. Linearization specifies the relative positions of the children of an AMR node to their mother and to each other. The focus of this paper is on the specific linearization submodule that deals with the problem of mapping thematic roles to sur-

face positions.

In (Dorr et al., 1998), the grammar formalism provided as part of the natural language realization engine Nitrogen (Langkilde and Knight, 1998a; Langkilde and Knight, 1998b) was used to implement a linearization grammar. The Nitrogen grammar formalism is unification based and it provides a small number of tools to recast and linearize AMRs. One of this formalism's limitations is the use of an interpreted grammar. Currently we are using a different linearization engine, Oxygen (Habash, 2000). Oxygen is an efficient language-independent linearization engine. Linearization grammars for Oxygen are written using oxyL, a powerful linearization grammar description language that has the power of a programming language with the focus on natural language linearization. oxyL grammars are compiled into programs that run independently.

## 4 The Thematic Hierarchy

The unordered nature of siblings under an AMR node complicates the mapping between AMR relations and their surface positions. In the case of thematic role ordering, the situation is more complicated by the lack of one-to-one mapping between a particular thematic role and an argument position. For example, a theme can be the subject in some cases and it can be the object in others or even an oblique. Observe *cookie* in (1).

(1)  (i)   John ate *a cookie* (object)

    (ii)  *the cookie* contains chocolate (subject)

    (iii) she nibbled *at a cookie* (oblique)

To solve this problem, a thematic hierarchy is used to determine the argument position of a thematic role based on its co-occurrence with other thematic roles. Several researchers have proposed different versions of thematic hierarchies (see (Jackendoff, 1972; Carrier-Duncan, 1985; Bresnan and Kanerva, 1989; Kiparsky, 1985; Larson, 1988; Giorgi, 1984; Wilkins, 1988; Nishgauchi, 1984; Alsina and Mchombo, 1993; Baker, 1989; Grimshaw and Mester, 1988)).[2] The hierarchy proposed in (Dorr et al., 1998) differs from these in that it separates (non-adjunct) arguments from obliques (i.e. adjunct arguments) and provides a more complete list of thematic roles (31 roles overall) than those of previous approaches (maximum of 8 roles). See Figure 3 for a listing of the most common thematic roles used. The following is the final thematic hierarchy for arguments.

(2)  ```special case -- ag src th
     ext > ag > instr > th > perc > ```*else*

In the case of the occurrence of theme alone, it is mapped to first argument position. If a theme and an agent occur, the agent is mapped to first argument position and the theme is mapped to second argument position. When an agent, a theme and a

---

[2]For an excellent overview and a comparison of different thematic hierarchies see (Levin and Rappaport Hovav, 1996).

Figure 5: Oxyl Thematic Hierarchy Implementation

```
:Recast &TH-order
(@this <? (:movsrc / (:src))
          (&and (&ex :ag) (&ex :th))
      <! ((:subj :obj1 :obj2) /
          (:ext :sub :ag :instr :movsrc :th :src :perc
           :goal :mod-poss :mod-loc :mod-pred :loc
           :poss :pred :prop :time :ben :purp)))
:Rule %S (-> (@subj @inst @obj1 @obj2))
```

source co-occur, The order in the hierarchy is violated as in $John_{ag}$ charged $Paul_{src}$ $\$40_{th}$. The term **ext** is used to handle verbs that violate the thematic hierarchy. It, **ext**, refers to an externally marked thematic role such as the perceived *John* in $John_{perc}$ pleases $Mary_{th}$ versus $Mary_{th}$ likes $John_{perc}$. This information is provided in the RLCS lexicon entry using the special marker :EXT. The use of the thematic hierarchy eliminates the need to specify the thematic role to surface position mapping in every verb lexicon entry.

As for the ordering of obliques, an ad hoc order was established:

(3)  `particle > mod-prop() > perc() > th() >`
     `purp() > mod-loc() > mod-pred() > src() >`
     `goal() > mod-poss() > ben()`

Note that the order of obliques is not a strict hierarchy but rather a possible topological sort. A more detailed discussion is available in (Dorr et al., 1998).

### 4.1 Thematic Hierarchy Implementation

Oxygen's linearization grammar description language, oxyL, provides a hierarchical data recasting operator that simplifies the implementation of thematic hierarchy mapping[3] (see Figure 5). The top part of Figure 5 defines the thematic hierarchy ordering as follows: given the current node (`@this`), conditionally recast (`<?`) the relation `:src:` into `:movsrc` if it co-occurs with `:ag` and `:th`; then hierarchically recast (`<!`) all available argument thematic roles into the grammatical roles `:subj`, `:obj1` and `:obj2`. The Linearization rule `%S` specifies the relative position of the arguments to the verb (`@inst`). The separation between Recasting and Linearization breaks up the problem of mapping a thematic role to a surface position into two subproblems: mapping a thematic role into a grammatical role (*subject,object*) and mapping a grammatical role into a surface position. The recasting and linearization rules are only fired if the AMR node being linearized is a verb. The linearization rule in our implementation specifies the relative location of the obliques. They are permuted at the end of the sentence.

---

[3] Another example of hierarchically ordered linguistic phenomena is the linearization of auxiliaries relative to the negative particle in the English verb phrase. The auxiliaries are strictly ordered by the part of speech (Modal Have Be+en Be+ing). The negative particle 'not' must appear after the first auxiliary regardless of its part of speech. A hierarchical mapping of the auxiliaries into (Aux1 Aux2 Aux3 and Aux4) is a simpler solution than listing all combinations.

### 4.2 Incorporation of Spanish

The linearization component that includes the thematic hierarchy mapping was implemented using Oxygen (Habash, 2000). The linearization grammar was very simple concentrating on argument word order relative to the verb using the same thematic hierarchy described in Figure 5.

To incorporate Spanish in our current implementation, we replaced complex Spanish morphology with the simple 'near-future' construction (*va ir a + INF*). For example, $alguien_{ag}$ va a colocar $algo_{th}$ en $algo_{goal}$. In addition to the lack of a complete phrase structure for parts of speech other than verbs, the Spanish linearization grammar doesn't handle Pro-drop or clitics. In principle, both phenomena can be handled with a recast rule that would fire after the thematic hierarchy recast. In the case of pro drop, it conjugates the verb and makes the subject null. And in the case of clitics, it adds a clitic that matches the gender and number of the object.

In the next section we evaluate the use of the thematic hierarchy for English and Spanish generation. The fact that English and Spanish are both SVO languages doesn't lessen the validity of the evaluation since the role of thematic hierarchies is not to map the thematic roles to surface positions but rather to the syntactic level (i.e. agent, theme, goal to subject, object and indirect object). Final linearization is responsible for placing the subject and object appropriately on the surface. The similarity of surface word order between Spanish and English should be seen as a normalization factor in testing the mapping from thematic roles to grammatical roles.

## 5 Evaluation

In this evaluation, a test corpus of 453 simple CLCSes corresponding to all Levin English verb classes and alternations was constructed semi-automatically. The test corpus size guarantees large-scale coverage over verb behavior and thematic role combinations, which is exhaustive for our purpose. The CLCSes were constructed by randomly selecting an LCS verb entry from each class from the English verb class and filling all its argument positions with simple noun phrases (e.g. $something_{th}$, $someone_{ag}$, etc.) or simple subordinate clauses (e.g. (to do something)$_{prop}$, (to be someproperty)$_{mod-prop}$, etc.). Table 1 shows some sample English sentences corresponding to the CLCSes in the test corpus.

For the purposes of this evaluation, statistical extraction was disabled because we do not have a Nitrogen bigram model for Spanish.

The CLCS test corpus was fed to the generation system in two different runs each of which using a different target language lexicon and oxyL linearization grammar. The results of the generation are passed to two speakers of English and Spanish respectively to evaluate the word order of the realized text. Evaluators were asked to mark sentences as being acceptable or not acceptable as far as the word-

| Verb Class | Example |
|---|---|
| 2 | something$_{ag}$ wanted something$_{th}$ (to do something$_{th}$)$_{prop}$ |
| 10.5 | someone$_{ag}$ stole something$_{th}$ from something$_{src}$ for something$_{ben}$ |
| 22.1.C | someone$_{ag}$ mixed something$_{th}$ into something$_{goal}$ |
| 29.1.B | someone$_{th}$ considered something$_{perc}$ (to be someproperty$_{pred}$)$_{mod-pred}$ |
| 45.2.A | someone$_{ag}$ folded something$_{th}$ with something$_{inst}$ |
| 55.1.C | someone$_{th}$ continued (to do something$_{th}$)$_{prop}$ |

Table 1: CLCS Test Corpus Examples

| $N = 453$ | Generated Classes | Word Order Error |
|---|---|---|
| English | 428 | 9% (40 classes) |
| Spanish | 254 | 2% (4 classes) |

Table 2: Initial Evaluation Results

Figure 6: New English Thematic Hierarchy

```
:Recast &TH-order
(@this <? (:mov / (:src :goal :ben))
          (&and (&ex :ag) (&ex :th))
       <! ((:subj :obj1 :obj2) /
          (:ext :sub :ag :instr :mov :th :src :perc
          :goal :mod-poss :mod-loc :mod-pred :loc
          :poss :pred :prop :time :ben :purp)))
:Rule %S (->(@subj @inst @obj1 @obj2))
```

order of the arguments relative to the verb[4]. Some of the English and Spanish sentences failed the lexical selection process due to problems with lexicon entries; these sentences never made it to linearization.

In the cases that survived, the lexical selection process appropriately generated multiple sentences for each CLCS. In the case of English, they all correctly corresponded to various related alternations of the main verb. For example, each of the two subclasses defining the dative alternations for the verb *send* generated each other (i.e. *John sent a book to Paul* and *John sent Paul a book*). There were also cases of overgeneration resulting from preposition under-specification, which is inconsequential to our evaluation(e.g. *go (to,toward,towards,to at,etc.) somewhere*).

On the other hand, in Spanish, there were many more sentences that should not have been generated. In theory, the lexical selection process limits the number of choices using the LCS entry of the Spanish verbs. But that process is only as good as the lexicon entries are. In cases where a bad sense is allowed in the translation, the sentence involved is dropped from the evaluation. This evaluation was quite helpful in pinpointing the locations of problems in our Spanish (and also English) lexicons. Table 2 displays the results of the evaluation. The first column represents the number of generated classes or CLCS instances (out of N = 453) that actually went through the whole system. Most failures in Spanish generation are due to missing verb entries (29% of all input classes). An additional 5% of classes was dropped out of the evaluation for having no correct sense output. The second column describes the ratio of classes with partially wrong or fully wrong word order output to the number of generated classes. In English, out of 428 classes, 30 classes had partially wrong output and 10 classes had no correct output. In Spanish, out of the 254 classes that generated output, only four classes had wrong word order output.

The next section describes the errors encountered in the evaluation and how they were fixed.

---

[4] Actually, the evaluation contained several other criteria that are more relevant to evaluating lexical selection such as completeness of argument realization and appropriateness of sense selection.

## 6 Discussion

The word-order errors in the English test belong to one of two types: First, there are lexicon errors where specific realization information such as :EXT is missing from an entry. This problem appeared in three subclasses of class 41.3.1 (Simple Verbs of Dressing: don, doff and wear). In our lexicon, *clothes*, the object for all three verbs, is considered the theme and the subject of the sentence is the goal, source and location respectively. Fixing these cases is a matter of adding the appropriate piece of information in the lexicon. The second type of errors were true thematic hierarchy errors: The case of agent-benefactor-theme co-occurrence such as *John bought Paul a house* and agent-goal-theme co-occurrence such as *John gave Paul a house*. These two should be part of the special case of the thematic hierarchy that deals with English verbs' indirect objects. Figure 6 displays the updated thematic hierarchy for English. In this implementation, a temporary role :MOV is created to mark source, goal or benefactor as *moved* arguments in a special conditional recasting step that depends on the co-occurrence of any of these roles with an agent and a theme.

The Spanish errors are much less than the English and are basically a subset of the first type of errors described above. The fact that the special case of the thematic hierarchy for English was included and it did not cause any problem to the Spanish is not surprising since Spanish lexical selection doesn't allow the thematic roles agent and theme to co-occur with the arguments source, goal or benefactor. The third argument is always generated as an oblique. For example, *Juan le dio un libro a Paolo* and *Juan le compró un libro a Paolo*[5]. The updated thematic hierarchy for Spanish is described in Figure 7.

We ran a second evaluation that uses the new implementations. The results are presented in Table

---

[5] We are aware that a more fluent Spanish would move the oblique (*a Paolo*) closer two the verb as in *Juan le dio a Paolo un libro* and *Juan le compró a Paolo un libro*. However this is not part of the focus of our evaluation. The behavior of obliques is something we plan to investigate in a separate study.

Figure 7: New Spanish Thematic Hierarchy

```
:Recast &TH-order
(@this <! ((:subj :obj1 :obj2) /
           (:ext :sub :ag :instr :th :src :perc
            :goal :mod-poss :mod-loc :mod-pred :loc
            :poss :pred :prop :time :ben :purp)))
:Rule %S (->(@subj @inst @obj1 @obj2))
```

| N = 453 | Generated Classes | Word Order Error |
|---|---|---|
| English | 428 | 1% (5 classes) |
| Spanish | 254 | 2% (4 classes) |

Table 3: Final Evaluation Results

3. For English, all of the classes with partially correct word order and half of the incorrect word order classes were corrected (88% of all erroneous classes). In the Spanish case, as expected, the results did not change.

Clearly, the results show that the use of a thematic hierarchy for generating both English and Spanish word order is successful and is supportive of earlier work (Dorr et al., 1998). The next step in this ongoing investigation is to test the use of the thematic hierarchy on a language with a different word order from English.

## 7 Future Work

A major remaining step is to correct the problems in the English and Spanish lexicons and to investigate the source of errors and incorrect sense selection. An investigation in the behavior of obliques in Spanish is necessary to produce fully fluent Spanish output. We have started investigating Chinese generation using thematic hierarchies; a preliminary study has shown promising results.

## 8 Acknowledgments

## References

Alsina, A. and Mchombo, S. (1993). Object Asymmetries and the Chichewa Applicative Construction. In Mchombo, S., editor, *Aspects of Automated Natural Language Generation*, pages 1–46. CSLI Publications, Center for the Study of Language and Information, Stanford, CA.

Baker, C. (1989). *English Syntax*. The MIT Press, Cambridge, MA.

Bresnan, J. and Kanerva, J. (1989). Locative Inversion in Chichewa: A Case Study of Factorization in Grammar. *Linguistic Inquiry*, 20:1–50.

Carrier-Duncan, J. (1985). Linking of Thematic Roles in Derivational Word Formation. *Linguistic Inquiry*, 16:1–34.

Dorr, B. J. (1997). Large-Scale Acquisition of LCS-Based Lexicons for Foreign Language Tutoring. In *Proceedings of the ACL Fifth Conference on Applied Natural Language Processing (ANLP)*, pages 139–146, Washington, DC.

Dorr, B. J., Habash, N., and Traum, D. (1998). A Thematic Hierarchy for Efficient Generation from Lexical-Conceptual Structure. In *Proceedings of the Third Conference of the Association for Machine Translation in the Americas, AMTA-98, in Lecture Notes in Artificial Intelligence, 1529*, pages 333–343, Langhorne, PA.

Dorr, B. J., Hendler, J., Blanksteen, S., and Migdaloff, B. (1993). Use of Lexical Conceptual Structure for Intelligent Tutoring. Technical Report UMIACS TR 93-108, CS TR 3161, University of Maryland.

Dorr, B. J. and Olsen, M. B. (1996). Multilingual Generation: The Role of Telicity in Lexical Choice and Syntactic Realization. *Machine Translation*, 11(1–3):37–74.

Giorgi, A. (1984). Toward a Theory of Long Distance Anaphors: A GB Approach. *The Linguistic Review*, 3:307–361.

Grimshaw, J. and Mester, A. (1988). Light Verbs and Theta-Marking. *Linguistic Inquiry*, 19:205–232.

Habash, N. (2000). oxyGen: A Language Independent Linearization Engine. In *Fourth Conference of the Association for Machine Translation in the Americas, AMTA-2000*, Cuernavaca, Mexico.

Jackendoff, R. (1972). Grammatical Relations and Functional Structure. In *Semantic Interpretation in Generative Grammar*. The MIT Press, Cambridge, MA.

Jackendoff, R. (1983). *Semantics and Cognition*. The MIT Press, Cambridge, MA.

Jackendoff, R. (1990). *Semantic Structures*. The MIT Press, Cambridge, MA.

Jackendoff, R. (1996). The Proper Treatment of Measuring Out, Telicity, and Perhaps Even Quantification in English. *Natural Language and Linguistic Theory*, 14:305–354.

Kiparsky, P. (1985). Morphology and Grammatical Relations. unpublished ms., Stanford University.

Langkilde, I. and Knight, K. (1998a). Generation that Exploits Corpus-Based Statistical Knowledge. In *Proceedings of COLING-ACL '98*, pages 704–710.

Langkilde, I. and Knight, K. (1998b). The Practical Value of N-Grams in Generation. In *International Natural Language Generation Workshop*.

Larson, R. (1988). On the double object construction. *Linguistic Inquiry*, 19:335–391.

Levin, B. and Rappaport Hovav, M. (1996). From Lexical Semantics to Argument Realization. Technical report, Northwestern University.

Levow, G.-A., Dorr, B., and Katsova, M. (2000). Construction of Chinese-English Semantic Hierarchy for Cross-Language Retrieval. In *Proceedings of the Workshop on English-Chinese Cross Language Information Retrieval, International Conference on Chinese Language Computing*, Chicago, IL.

Nishgauchi, T. (1984). Control and the Thematic Domain. *Language*, 60:215–260.

Traum, D. and Habash, N. (2000). Generation from Lexical Conceptual Structures. In *Proceedings of the Workshop on Applied Interlinguas, North American Association of Computational Linguistics/Applied Natural Language Processing Conference, NAACL/ANLP-2000*, Seattle, WA.

Wilkins, W., editor (1988). *Syntax and Semantics 21: Thematic Relations*. Academic Press, San Diego, CA.