DEVELOPMENT OF INTERLINGUAL LEXICAL CONCEPTUAL STRUCTURES WITH SYNTACTIC MARKERS FOR MACHINE TRANSLATION

UMIACS TR 95-16, CS TR 3412

by

Bonnie Dorr, Jye-hoon Lee, Clare Voss, and Sungki Suh
Institute for Advanced Computer Studies
and Department of Computer Science
University of Maryland
A.V. Williams Building
College Park, MD 20742

for

U.S. Army Research Office
Center for Command, Control, and Communications Systems (C3)
(Mr. George Yaeger)

February, 1995

Contract No. DAAL03-91-C-0034
TCN Number 93005
Scientific Services Program

Abstract

This document reports on research conducted at the University of Maryland for the Korean/English Machine Translation (MT) project. Our primary objective was to develop an interlingual representation based on lexical conceptual structure (LCS) and to examine the relation between this representation and a set of linguistically motivated semantic classes. We view the work of the past year as a critical step toward achieving our goal of building a generator: the classification of LCS's into a semantic hierarchy provides a systematic mapping between semantic knowledge about verbs and their surface syntactic structures.

We have focused on several areas in support of our objectives: (1) investigation of morphological structure including distinctions between Korean and English; (2) porting a fast, message-passing parser to Korean (and to the IBM PC); (3) study of free word order and development of the associated processing algorithm; (4) investigation of the aspectual dimension as it impacts morphology, syntax, and lexical semantics; (5) investigation of the relation between semantic classes and syntactic structure; (6) development of theta-role and lexical-semantic templates through lexical acquisition techniques; (7) definition a mapping between KR concepts and interlingual representations; (8) formalization of the lexical conceptual structure

Foreword

This document reports on research conducted at the University of Maryland for the Korean/English Machine Translation (MT) project. This effort has involved Dr. Bonnie Dorr, professor in computational linguistics, and three graduate students in computer science. Biographical sketches are given below. In addition, the Maryland team is in close collaboration with Martha Palmer at the University of Pennsylvania and Patrick Saint-Dizier at IRIT in Toulouse, France.

Bonnie Dorr received her bachelor's degree, summa cum laude, in computer science from Boston University, Boston, MA, in 1984, and her S.M. and Ph.D. in computer science from the Massachusetts Institute of Technology, Cambridge, MA, 1990. Between September 1990 and June 1992, she was a research associate with the Institute for Advanced Computer Studies at the University of Maryland, College Park; since then she has been an assistant professor in the Computer Science Department. She has over a decade of experience with natural language processing systems, has been a part of the natural language processing group at the University of Maryland since 1990, and has undertaken extensive practical industry consulting as well as work with multilingual translation development and evaluation at MITRE. Dr. Dorr has recently received the Alfred P. Sloan Research Fellowship and a NSF Young Investigator Award for her work on Large-Scale Interlingual Machine Translation. She is the author of several published papers on natural language processing as well as a book on the lexicon in machine translation. She is a member of ACL, ACM, AMTA, AAAI, Phi Beta Kappa, and Pi Mu Epsilon, and Sigma Xi.

Jye-hoon Lee received his bachelor's degree in mechanical engineering at Seoul National University, Seoul, Korea in 1989, studied Computer Science at Iowa State University between 1989 and 1991 and then became a doctoral student in the computer science graduate program at the University of Maryland. Since January 1993, Mr. Lee has been involved in the use of corpora for the acquisition of lexical-semantic features associated with verb meaning in the context of machine translation. Over the last year, he has also written a transliteration program from Yale romanization to Korean (and vice versa) and has analyzed problems that arise in translating between Korean and English. Mr. Lee expects to graduate in 1996.

Sungki Suh received his bachelor's degree in English language and literature at Seoul National University, Seoul, Korea in 1986 and his M.A. in linguistics at State University of New York at Buffalo, New York. He received his Ph.D. in Linguistics, graduating 3 months prior to the completion of the project in 1994. He is currently a professor in the Language Research Institute at Seoul National University. Dr. Suh specialized in the area of computational linguistics and is one of the few linguists in the world whose expertise is in Korean syntax. He has recently published papers in prestigious linguistics conferences and has extensive research experience in the linguistics department at Maryland.

Clare Voss received her bachelor's degree in linguistics at the University of Michigan, Ann Arbor, MI, her M.A. in psychology at the University of Pennsylvania, Philadelphia, PA, and is currently a doctoral student in the computer science department at the University of Maryland. Ms. Voss has been a researcher in computational linguistics for over seven years and has specialized in the cross-linguistic study of diverse languages such as French, German, Portuguese, and Estonian. She is currently working on a lexicon-based machine translation system called LEXITRAN under Dr. Dorr, which has a comprehensive coverage of spatial and aspectual relations and uses the principles-and-parameters framework assumed for the current project. She is also designing a MT developer's tool called ILustrate which allows computational linguists to test different version of the IL representation. Ms. Voss expects to receive her Ph.D. in fall of 1996.

Martha Palmer is an adjunct professor in the Institute for Research in Cognitive Science at the University of Pennsylvania. Beginning as a Sloan Fellow at the University of Pennsylvania in 1981, and continuing as a Senior Research Scientist at Unisys in the 80's, Dr. Palmer has investigated the formalizations of underlying conceptual representations of verbs in several different types of subdomains, including physics word problems and a natural language interface to a graphics system. During a 3-year stay in Singapore, Dr. Palmer extended her methods of verb representation to French and Chinese, and investigated their utility for resolving cross-linguistic divergences. More recently, she has been collaborating with Dr. Dorr and Dr. Saint-Dizier on conceptual representations of English and French verbs that are necessary for accurate lexical selection in machine translation. She has also investigated the application of lexicalized grammars to the syntactic and semantic structure of Korean sentences.

Patrick Saint-Dizier has directed work in natural language at the Universite Paul Sabatier in Toulouse, France since 1986, laying the foundations for addressing problems in understanding and generation of language. After collecting real world data and discovering that certain linguistic constructs and lexical items recur regularly in different languages, Dr. Saint-Dizier's group has pursued a formal characterization of the linguistic theory of the lexicon and has developed constraint logic programming based techniques for translating natural language based on this formalism.

Acknowledgments

This work was supported by the Center for Command, Control, and Communications Systems (C3) (Mr. George Yaeger) under the auspices of the U.S. Army Research Office Scientific Services Program administered by Battelle (Delivery Order 584, Contract No. DAAL03-91-C-0034).

Contents

1 Introduction **5**

2 Transliteration and Morphology **5**
2.1 Hangul Converter 5
2.2 Irregular Verbs 5
2.3 Nominal Suffix 6
2.4 Verbal Endings 7
2.5 Korean Auxiliary Verbs 9

3 Message-Passing Parser and LCS Composition **10**
3.1 Message Passing Paradigm 11
3.2 Implementation of Principles 14
3.2.1 $\overline{\mathrm{X}}$ Theory 14
3.2.2 Trace Theory 14
3.2.3 Case Theory 14
3.3 Implementation of Parameters 15
3.3.1 $\overline{\mathrm{X}}$ Theory 16
3.3.2 Trace Theory 16
3.3.3 Case Theory 16
3.3.4 Summary of Parameter Settings for English and Korean 17
3.4 Results of Time Test Comparisons 17
3.5 Implications for Machine Translation 17
3.6 LCS Composition 20

4 Translation Issues **21**
4.1 Paraphrases and Resumptive Pronouns 21
4.2 Use of Chinese Characters for Semantic Distinctions 22
4.3 Syntactic Structure of Korean 23
4.4 Scrambling as Case-Driven Obligatory Movement 23
4.5 Empty Categories 24

5 Temporal/Aspectual Analysis **26**
5.1 Possible Tenses in English and Korean 26
5.2 The Place of Aspect in an IL-Based MT System 28
5.3 Lexical and Grammatical Aspect: Olsen 29

6 Lexicon Development and Semantic Classification of Verbs **30**
6.1 Semi-Automatic Tools for Lexicon Development 30
6.2 Automated Lexicon Construction: LEXICALL 34
6.3 Verbs of Motion 36
6.3.1 Compound Verbs Representing Spontaneous Motion 36
6.3.2 Verbs Representing Caused Motion 37
6.4 Lexical Semantics and French Data 37
6.4.1 Spatial predicates for Dorr's LCSs: English 38
6.4.2 Spatial predicates for Saint-Dizier's LCSs: French 39

7 Automatic Acquisition of Lexical Entries **39**
7.1 Construction of Thematic Grids from LDOCE codes 41
7.1.1 Assignment of LDOCE-Based Thematic Grids to English Verbs 41
7.2 Construction of Thematic Grids from Levin's Classification 46
7.2.1 Assignment of Levin-Based Thematic Grids to Lexical Entries 48
7.2.2 Results of Brute-Force Classification Strategy 50
7.2.3 Application of Linguistic Techniques for Levin-Based Classification 51

8 Interface Between LCS in KR **55**
8.1 A Brief Architectural Description 56
8.2 Lexical Mismatches 56
8.2.1 Space of Lexical Mismatches 58
8.2.2 Our Focus 60
8.3 Combining Aspects of Lexical Semantics and KB 60
8.3.1 Benefits of Retaining the LCS formalism 60
8.3.2 Encoding the LCS formalism in the KB 61
8.4 LOOM Implementation 62
8.4.1 Before Lexical Match Time 62
8.4.2 Lexical Selection and Mismatch Handling 64
8.4.3 Inner Workings of the Lexical Selection Algorithm 65
8.5 Areas for Future Investigation: KB/IL 66

9 Relation of LCS to TAGs **66**
9.1 Grammar for LCS 66
9.2 Tree Adjoining Grammar (TAG) 67
9.3 LCS Operation Requirements 69
9.4 Implications and Future Directions 72

A Hangul Conversion Program **78**

B BNF for Syntactic Output of Parser **78**

C Korean Lexical Entries for Parser and LCS Composition **81**
C.1 Parser Entries 81
C.2 LCS Entries 82

D Parser and LCS Composition Output **86**
D.1 E: John married Sally 86
D.2 E: John helped Bill 86
D.3 K: John-i Sally-wa kyelhonhayssta (John married (with) Sally) 87
D.4 K: Bill-un John eykeyse towum-ul patassta (Bill received help from John) 87
D.5 K: John-i Bill-eykey towum-ul cwuessta. (John gave help to Bill) 88
D.6 K: John-i Bill-ul towassta (John helped Bill) 89

E Korean Morphology Dictionary **90**

1 Introduction

Over the last year, the computational linguistics group at Maryland has been involved in the problem of interlingual machine translation of Korean and English. We have adopted a pivot form called Lexical Conceptual Structure (LCS), was developed on the basis of work by (Jackendoff, 1983), (Jackendoff, 1990). Our primary objective was to develop an interlingual representation based on lexical conceptual structure (LCS) and to examine the relation between this representation and a set of linguistically motivated semantic classes. We view the work of the past year as a critical step toward achieving our goal of building a generator: the classification of LCS's into a semantic hierarchy provides a systematic mapping between semantic knowledge about verbs and their surface syntactic structures.

We have focused on several areas in support of our objectives: (1) investigation of morphological structure including distinctions between Korean and English; (2) porting a fast, message-passing parser to Korean (and to the IBM PC); (3) study of free word order and development of the associated processing algorithm; (4) investigation of the aspectual dimension as it impacts morphology, syntax, and lexical semantics; (5) investigation of the relation between semantic classes and syntactic structure; (6) development of theta-role and lexical-semantic templates through lexical acquisition techniques; (7) definition a mapping between KR concepts and interlingual representations; (8) formalization of the lexical conceptual structure Each of these areas is addressed, in turn, in the following sections.

2 Transliteration and Morphology

We continued development on the Hangul Converter and also made some headway on developing rules that will serve as input to the Kimmo system for a morphological analyzer/synthesizer.

2.1 Hangul Converter

At the beginning of 1994, we finalized the program that allows us to convert between Hangul and romanized forms. Appendix A provides usage notes for this program. This year, we have added the ability to use Chinese characters. There exist more than 100,000 Chinese characters. Because there is no alphabet system in Chinese, this facility requires a 17-bit coding system to distinguish among all possibilities. ($2^17 = 131072$)

KSC-5601, which is a complete Hangul code, defines most commonly used 3000 Chinese characters. However, the KSC-5601 code was defined in Korea; thus, no other Chinese characters are allowed other than 3000 Chinese characters defined in KSC-5601.

A difficulty that we faced when we installed Chinese characters is that there is a many-many mapping between Chinese characters and Korean characters. Chinese characters have potentially more than one Korean sound, and Korean characters, which collectively have 2350 different sounds, can be mapped to many different Chinese characters. Thus we were forced to build, by hand, a very large Chinese character table. Our approach was to distinguish among several Chinese characters of the same sound using a signal for each case (e.g., Ka1, Ka2, ...).

We added an external romanization mapping table that allows the user to define mapping functions between a 2 byte character and a given romanization scheme. We added 3000 Chinese characters in the 5601 standard.

2.2 Irregular Verbs

We developed the following 6 rules of morphology for irregular verbs using (Ihm et al., 1988).

1. n

```
l --> 0 / __ +  s
                p
```

2.
```
p --> o / o __ + V(owel)
p --> wu / __ + V
Exceptions: nelp-ta, cop-ta, cip-ta, cap-ta, pwutcap-ta,
ip-ta, ep-ta, ppop-ta, ssip-ta
```

3.
```
t --> l / __ + V
Exceptions: mwut-ta, mit-ta, et-ta, ssot-ta, pat-ta, tat-ta
```

4.
```
u --> l / l __ +  a
                  e
```

5.
```
                           n
h --> 0 / #(C)V(C)CV __ +  l
                           m
                           s
```

6.
```
s --> 0 / __ + V
Exceptions: ppayas-ta, wus-ta, pes-ta, ssis-ta
```

2.3 Nominal Suffix

We undertook an analysis of the nominal suffix from (Ihm et al., 1988), which uncovered the following information:

I. Case Marker

- (1) Nominative case marker: ka/ i, kkeyse(-honorific)
  `ka --> i / C(onsonant) + ___`
- (2) Accusative case marker: lul/ ul
  `lul --> ul / C + ___`
- (3) Genitive case marker: uy

II. Postposition

- (1) Instrumental: lo/ ulo
  `lo --> ulo / C + ___ (Exception: ulo --> lo / l + ___ )`
- (2) Reason or Source: lo/ ulo, ey
- (3) Status: lo/ ulo
- (4) Resultative: lo/ ulo
- (5) Locative
  - (a) (default): ey
  - (b) [+Animate] (=Dative): ey-key, hanthey, kkey (-honorific)
  - (c) Source: ey-se, ey-key-se, hanthey-se, pwuthe (-beginning point)
  - *(d) Direction: lo/ ulo
  - (e) Ending point: kkaci
  - (f) Eventive: eyse
- *(6) Temporal

   (a) Eventive: ey
   (b) Beginning point: eyse, pwuthe
   (c) Ending point: kkaci

(7) Measure: ey
(8) Commitative: wa/ kwa, lang/ ilang, hako
    `wa --> kwa / C + ___`
    `lang --> ilang / C + ___`

III. Delimiter

(1) Topic: nun/ un
    `nun --> un / C + ___`
    cf: Topic marking of the subject is required when the sentence is classified as a depictive statement.
(2) Only: man
(3) Too: to
(4) Even: cocha, mace
(5) Each: mata
(6) Unselective or Emphasis: na/ ina
    `na --> ina / C + ___`
(7) Amount: (Only in interrogative sentences): na/ ina

** When a delimiter is attached to the subject/ object NP, nominative/ accusative case marker on the NP must be deleted. (Case III.(2) above is an exception.)

```
• John-un /*John-i-un    Mary-to/ *Mary-ul-to    coahanta
      -Top     -Nom-Top     -too      -Acc-too     like
  'John likes Mary, too.'
  ('John, even Mary likes him.')
```

** Delimiters can be attached to adverbs/ verbs as well as nouns.

IV. Others

(1) Comparative: pota, kathi, chelem
(2) Coordinate conjunction
   (a) And: wa/ kwa, lang/ ilang, hako
   (b) Or: na/ ina

2.4 Verbal Endings

We undertook an analysis of Verbal Endings from (Ihm et al., 1988). In the description below, unless specified otherwise, the initial vowel of the verbal ending (= 'u' or 'e') is deleted when the verbal stem ends with a vowel. (Verbs whose stem ends with '-o'/'-wu' sound are exceptions to this generalization.) Also, the initial vowel 'e' of verbal ending changes to 'a' when the last syllable of verbal stem includes '-a' or '-o' sound.

I. Terminative Endings

*Terminative endings represent the mood type of sentences. They are classified based on the speech level. Speech level is mainly determined by the hearer's age, social status etc. relative to the speaker's.

(1) Declarative
  (a) [super high]: 'upnita'
  (b) [high]: 'eyo', 'ciyo'
  (c) [mid-low]: 'so', 'ne'
  (d) [low]: 'ta', 'e'

(2) Interrogative
  (a) [super high]: 'upnikka'
  (b) [high]: 'eyo', 'nayo'
  (c) [mid-low]: 'swu', 'na'
  (d) [low]: 'ni', 'nya'

(3) Imperative
  (a) [super high]: 'useyyo'
  (b) [high]: 'eyo'
  (c) [mid-low]: 'key'
  (d) [low]: 'ela', 'e'

(4) Propositive
  (a) [super high]: 'siciyo'
  (b) [high]: 'upsita', 'ciyo'
  (c) [mid-low]: 'use'
  (d) [low]: 'ca'

II. Adnominal Endings (-involving either a relative clause or a complex NP clause)

(1) Present tense
  (a) 'un': for adjectival verbs
  (b) 'nun': otherwise

```
• kem-un   koyangi
  black-Adnm   cat
  'A cat which is black'/'a black cat'

• chayk-ul   ilk-nun  salam
   book-Acc read-Adnm  man
  'A man who is reading a book'
```

(2) Future tense: 'ul'

(3) Past tense
  (a) (ess)'ten': implying reminiscence
  (b) 'un': used only with non-adjectival verbs

III. Adverbial Endings

(1) reason or cause: 'se' 'nikka' `--->`for, as

(2) weak contrast: 'nuntey' `--->`while

(3) conditional: 'umyen' 'ketun' `--->`if, when

(4) purpose: 'lyeko' 'le' `--->`in order to (do), for (do)ing

(5) prerequisite:'(e/a)ya' '(e/a)yaman' `--->`only when, only if

(6) goal: 'tolok' `--->`so that (one) may/can (do)

(7) concurrence: 'umyense' 'umye'

(8) contrast: 'ciman' `--->`although (cf: coordinate conjunction 'ciman')

(9) separate action: 'ta(ka)'

(10) greater degree:'ulswulok' `--->`the more (....the more)

(11) immediate sequence: 'ca' 'camaca' `--->` as soon as

IV. Nominal Endings -'um' and 'ki'

(1) [+tense]: 'um'

(2) [-tense]: 'ki'

*'um' must be accompanied by a case marker.

*'um' generally occurs with factive predicate, and 'ki' occurs with nonfactive predicate.

V. Coordinate Conjunction

(1) and: 'ko', 'se'
*'se' is used when the first conjunct precedes the second one in time sequence or when the first conjunct is subordinate to the second one. (cf:'kose')

(2) or: 'kena'

(3) but: 'ciman', 'una'

VI. (Quotative) Complementizer: 'ko'

*'ko' must be preceded by terminative endings

2.5 Korean Auxiliary Verbs

Our analysis of Korean auxiliary verbs from (Ihm et al., 1988) revealed that such verbs are classified primarily into two groups: One corresponds to an aspectual specification, and the other corresponds to the representation of a state which is different from the present.

I. Aspectual Specification

*V = main verb stem

(1a) V-e 'peli-ta': completion

(1b) V-e 'nay-ta': accomplishment

(1c) V-ko 'mal-ta': perfective(?)(-something is done at last)

(2a) V-e 'noh-ta': completion + duration

(2b) V-e 'twu-ta': duration

(2c) V-e 'kaci-ko': duration(?) (-must be used in the form of 'conjunction')

(3a) V-kon ‘ha-ta’: habitual

(3b) V-e ‘tay-ta’: repetition(?)

(4) V-ko ‘iss-ta’: progressive

II. The State differing from the present

(1) V-ko ‘siph-ta’: hope (-want/hope to V)

(2) V-na ‘siph-ta’: (speaker’s) guess

(3) V-nunka ‘ha-ta’: (speaker’s) guess
V-na ‘ha-ta’

(4) V-un ‘tus-siph-ta’: (speaker’s) expect/ guess
V-ul ‘tus-siph-ta’

(5) V-un ‘tus-ha-ta’: (speaker’s) expect/ guess
V-ul ‘tus-ha-ta’

(6) V-eya ‘ha-ta’: obligation (-have to V)

(7) V-un ‘cheyha-ta’: pretense (-pretend to V)

(8) V-ul ‘ppenha-ta’: almost (-almost did something, but no success/ completion) ---> past tense required

(9) V-ulye(ko) ‘ha-ta’: is about to V

(10) V-koca ‘ha-ta’: volition

(11) V-ulkka ‘ha-ta’: plan(-not decisive)

(12) V-ul ‘manha-ta’: is worthwhile to V

(13) V-na ‘po-ta’: (speaker’s) guess --->tense marker is not allowed

III. Others

(1) V-e ‘cwu-ta’: of benefit (-did something for others)

(2) V-e ‘po-ta’: trial

(3) V-ci ‘anh-ta’: negation (-does not V)

3 Message-Passing Parser and LCS Composition

Another area that we have investigated over the past year is that of efficient cross-linguistic parsing. The model we have adopted is a Message-Passing approach, based on Government-Binding (GB) Theory (Chomsky, 1981), (Haegeman, 1991), (van Riemsdijk and Williams, 1986). One of the drawbacks to alternative GB-based parsing approaches is that they generally adopt a filter-based paradigm. These approaches typically generate all possible candidate structures of the sentence that satisfy $\overline{\mathrm{X}}$ theory, and then subsequently apply filters in order to eliminate those structures that violate GB principles. (See, for example, (Abney, 1989), (Correa, 1991), (Dorr, 1991), (Fong, 1991), (Frank, 1990).) The current approach provides an alternative to filter-based designs which avoids these difficulties by applying principles to *descriptions* of structures without actually building the structures themselves. Our approach is similar to that of (Lin, 1993), (Lin and Goebel, 1993) in that structure building is deferred until the descriptions satisfy all principles; however, the current approach differs in that it provides a parameterization mechanism along the lines of (Dorr, 1993a) that allows the system to be ported to languages other than English. We focus particularly on the problem of processing head-final languages such as Korean.

We are currently incorporating the parser into our machine translation (MT) system, which we have now named PRINCITRAN.[1] In general, parsers of existing principle-based interlingual MT systems are exceedingly inefficient since they tend to adopt the filter-based paradigm. We combine the benefits of the message-passing paradigm with the benefits of the parameterized approach to build a more efficient, but easily extensible system, that will ultimately be used for MT. The algorithm has been implemented in C++ and successfully tested on well-known, translationally divergent sentences.

It is important to point out that the efficiency of the system is not simply a side effect of using an efficient programming language (i.e., C++), but that the algorithm is inherently efficient, independent of the programming language used for the implementation. The key to the efficiency of the algorithm is that it is based on a system that is already provably fast (Lin and Goebel, 1993). A formal, worst-case analysis reveals that the original CFG parsing algorithm is $O(|G|n^3)$ (Lin and Goebel, 1993), where n is the length of the input sentence and $|G|$ is the size of the grammar (the number of occurrences of non-terminal symbols in the set of grammar rules). This is an improvement of a factor of $|G|$ over standard CFG algorithms such as the Earley parser (Barton et al., 1987). This complexity measure is based on a serial simulation of parallel distributed message passing.

The extended version of the system differs only in that it includes attribute values, constraints, and movements, thus allowing a wider range of phenomena to be handled while still retaining the efficiency of the original algorithm. The complexity of this version has not yet been formally determined; however, in this article we will provide an intuitive sense of the inherent speed of this version by describing a set of experiments that show that the average time does not grow exponentially to the sentence length.[2]

The next section presents a general framework for parsing by message passing. The following section describes our implementation of GB principles as attribute-value constraints in the message-passing framework. We then present the parameterization framework, demonstrating the feasibility of handling cross-linguistic variation within the message-passing framework. A technique for automatic precompilation of parameter settings is then described. In the section on Time Test Comparisons, we compare the efficiency of the parser to that of the original CFG algorithm (Lin and Goebel, 1993) as well as Tomita's algorithm (Tomita, 1986) on a test suite of representative sentences. We conclude with a discussion about the implications of this approach and its applicability to the problem of multi-language MT; preliminary results on translationally divergent sentences in Korean and English are presented.

3.1 Message Passing Paradigm

There has been a great deal of interest in exploring new paradigms of parsing, especially non-traditional parallel architectures for natural language processing (Cottrell, 1989), (Waltz and Pollack, 1985), (Small, 1981), (Selman and Hirst, 1985), (Abney and Cole, 1985), (Jones, 1987), (Kempen and Vosse, 1989) (among many others). Recent work (Stevenson, 1994) provides a survey of symbolic, non-symbolic, and hybrid approaches. Stevenson's model comes the closest in design to the current principle-based message-passing model in that it uses distributed message passing as the basic underlying mechanism and it encodes GB principles directly (i.e., there are precise correspondences between functional components and linguistic principles). However, the fundamental goals of the two approaches are different: Stevenson's objective concerns the modeling of human processing behavior and producing a single parse at the end. Her system incorporates

[1] The name PRINCITRAN is derived from the names of two systems, UNITRAN (Dorr, 1993b) and PRINCIPAR (Lin, 1993).

[2] Presumably, runtime differences that are due solely to differences in efficiency of programming language will result in a curve that is similar to others, but at a different scale. We will show that, in fact, this is the result we obtained.

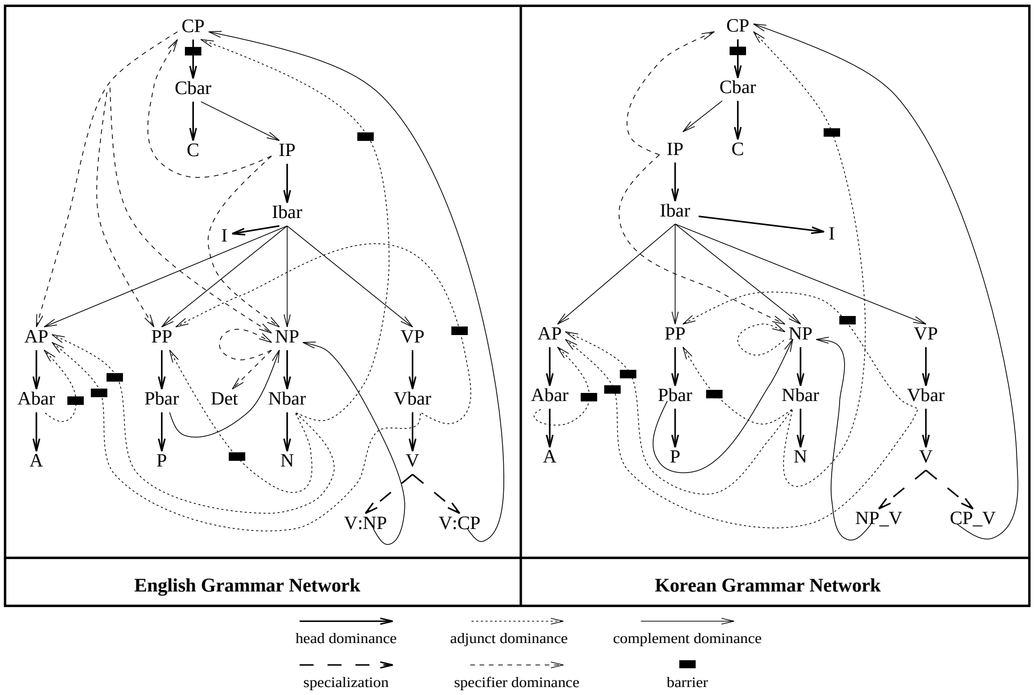

Figure 1: Network Representation of English and Korean Grammar

a number of psycholinguistic-based processing mechanisms for handling ambiguity and making attachment decisions.[3] Our model, on the other hand, is more concerned with efficiency issues, broad scale coverage, and cross-linguistic applicability; we produce all possible parse alternatives wherever disambiguation requires extra-sentential information.

We view Stevenson's approach as a special case of our own in that it would be possible to program our system so that it provides identical output for the same phenomena. The benefit of Stevenson's approach is that the linguistic constraints fall out from the machinery whereas our own design requires the constraints to be explicitly encoded in the processing data structures. On the other hand, a critical deficiency of Stevenson's design is that it is incapable of handling head-final languages unless the machinery undergoes a massive modification.

Our approach provides a language-independent processing mechanism that accommodates structurally different languages (e.g., head-initial vs. head-final) with equally efficient run times and it does not require major modifications to the underlying processing mechanism as each language is added. The approach extends the GB-based message-passing design proposed by (Lin, 1993), (Lin and Goebel, 1993). The grammar for each language is encoded as a network of nodes that represent grammatical categories (e.g., NP, Nbar, N) or subcategories, such as V:NP (i.e., a transitive verb that takes an NP as complement). Figure 1 depicts portions of the grammar networks used for English and Korean.

There are two types of links in the network: **subsumption links** (e.g., V to V:NP) and **dominance links** (e.g., Nbar to N). A dominance link from α to β is associated with an integer id that determines the linear order between β and other categories immediately dominated by α, and a binary attribute to specify whether β is optional or obligatory.[4]

[3] Stevenson's system currently handles structural ambiguity, but not lexical ambiguity; her research has focused almost entirely on the question of attachment decisions rather than on issues concerning selection of lexical items.

[4] For the purpose of readability, we have omitted integer id's in the graphical representation of the grammar network. Linear ordering is indicated by the starting points of links. For example, C precedes IP in the English network of figure 1.

Input sentences are parsed by passing messages in the grammar network. The nodes in the network are computing agents that communicate with each other by sending messages in the reverse direction of the links. Each node locally stores a set of items. An item is a triplet that represents a $\overline{\mathrm{X}}$ structure α: <surface-string, attribute-values, source-messages>, where *surface-string* is an integer interval [i,j] denoting the i'th to j'th word in the input sentence; *attribute-values* specifies syntactic features of the root node (β); and *source-messages* is a set of messages that represent immediate constituents of β and from which this item is combined. Each node has a completion predicate that determines whether an item at the node is "complete," in which case the item is sent as a message to other nodes.

When a node receives an item, it attempts to form new items by combining it with items from other nodes. Two items $<[i_1,j_1], A_1, S_1>$ and $<[i_2,j_2], A_2, S_2>$ can be combined if:

1. their surface strings are adjacent to each other: $i_2 = j_1+1$.
2. their attribute values A_1 and A_2 are unifiable.
3. the source messages come via different links: $\text{links}(S_1) \cap \text{links}(S_2) = \emptyset$, where links(S) is a function that, given a set of messages, returns the set of links via which the messages arrived.

The result of the combination is a new item:

$$<[i_1,j_2], \text{unify}(A_1, A_2), S_1 \cup S_2>.$$

The new item represents a larger $\overline{\mathrm{X}}$ structure resulting from the combination of the two smaller ones. If the new item satisfies local constraints (to be described in the next section), it is considered valid and saved in the local memory. Otherwise, it is discarded. A valid item satisfying the completion predicate of the node is sent further as a message to other nodes.

More specifically, the parsing algorithm consists of the following steps:

Step 1: Lexical Look-up Retrieve the lexical entries for all the words in the sentence and create a lexical item for each word sense. A lexical item is a triple: $<[i,j], av_{self}, av_{comp}>$, where [i,j] is an interval denoting the position of the word in the sentence; av_{self} is the attribute values of the word sense; and av_{comp} is the attribute values of the complements of the word sense.

Step 2: Message Passing For each lexical item $<[i,j], av_{self}, av_{comp}>$, create an initial message $<[i,j], av_{self}, \emptyset>$ and send this message to the grammar network node that represents the category or subcategory of the word sense. When the node receives the initial message, it may forward the message to other nodes or it may combine the message with other messages and send the resulting combination to other nodes. This initiates a message passing process which stops when there are no more messages to be passed around. At that point, the initial message for the next lexical item is fed into the network.

Step 3: Build a Shared Parse Forest When all lexical items have been processed, build a shared parse forest for the input sentence by tracing the origins of the messages at the highest node (CP or IP), whose `surface-string` component is the whole sentence. The links in the parse forest are derived from the grammar network links that are traversed during the tracing process. The structure of the parse forest is similar to (Tomita, 1986) and (Billot and Lang, 1989), but extended to include attribute values.

The parse trees representing the input sentence are retrieved from the parse forest one by one. The next section explains how constraints attached to the nodes and links in the network ensure that the parse trees satisfy all GB principles.

3.2 Implementation of Principles

GB principles are implemented as **local constraints** attached to nodes and **percolation constraints** attached to links. All items at a node must satisfy the node's local constraint. A message can be sent across a link only if it satisfies the link's percolation constraint. The idea of constraint application through feature passing among nodes is analogous to techniques applied in the TINA spoken language system (Seneff, 1992) except that, in our design, the grammar network is a *static* data structure; it is not *dynamically* modified during the parsing process. Thus, we achieve a reduction space requirements. Moreover, our design achieves a reduction in time requirements because we do not retrieve a structure until the resulting parse descriptions satisfy all the network constraints.

We will discuss three examples to illustrate the general idea of how GB principles are interpreted as local and percolation constraints. See (Lin, 1993) for more details.

3.2.1 $\overline{\text{X}}$ Theory

The central idea behind $\overline{\text{X}}$ theory is that a phrasal constituent has a layered structure. Every phrasal constituent is considered to have a head ($\text{X}^0 \equiv \text{X}$), which determines the properties of the phrase containing it. A phrase potentially contains a complement, resulting in a one-bar level ($\overline{\text{X}}$ $\equiv$ Xbar) projection; it may also contain a specifier (or modifier), resulting in a double-bar level ($\overline{\overline{\text{X}}}$ $\equiv$ XP) projection. The phrasal representation assumed in the current framework is the following:

(1) $[_{\text{XP}}$ Specifier $[_{\text{Xbar}}$ Complement X]]

We implement the relative positioning of Specifier, Complement, and Head constituents by means of dominance links as shown in each of the networks of Figure 1. In addition, adjuncts are associated with the Xbar level by means of an adjunct-dominance link in the grammar network.

The structure in (1) represents the relative order observed in Korean. We will see shortly that the linear order between the head X and the Complement is determined by a parameter of variation.

3.2.2 Trace Theory

A trace represents a position from which some element has been extracted.[5] The main constraint of Trace Theory is the Subjacency Condition, which prohibits movement across "too many" barriers. (The notion of "too many" is specified on a per-language basis as we will see shortly.)

An attribute named `barrier` is used to implement this principle. A message containing the attribute value `-barrier` is used to represent an $\overline{\text{X}}$ structure containing a position out of which a wh-constituent has moved, but without yet crossing a barrier. The value `+barrier` means that the movement has already crossed one barrier. Certain dominance links in the network are designated as barrier links (indicated in Figure 1 by solid rectangles). The Subjacency condition is implemented by the percolation constraints attached to the barrier links, which blocks any message with `+barrier` and changes `-barrier` to `+barrier` (i.e., it allows the message to pass through).

3.2.3 Case Theory

Case theory requires that every NP be assigned abstract case. The Case Filter (Chomsky, 1981) rules out sentences containing an NP with no case. Case is assigned structurally to a syntactic position governed by a case assigner. Roughly, a preposition assigns Oblique Case to a prepositional object NP; a transitive verb assigns Accusative Case to a direct object NP; and tensed Infl(ection) assigns Nominative Case to a subject NP.

[5]A trace is represented as t_i, where i is a unique index referring to an antecedent.

Feature	Significance
`+ca`	the head is a case assigner
`-ca`	the head is not a case assigner
`+govern`	the head is a governor
`-govern`	the head is not a governor
`-cm`	an NP m-commanded by the head needs case marking

Figure 2: Case Theory Attribute Values

Node	Local Constraint
P	assign `+ca` to every item
V	assign `+ca` to items with `-passive`
I	assign `+ca` to items with `tense` attribute

Figure 3: Constraints on Case Assignment

The implementation of case theory in our system is based on the following attribute values: `ca, govern, cm` (see Figure 2). The attribute values `+ca` and `+govern` are assigned by local constraints to items representing phrases whose heads are case assigners (e.g., tensed I) and governors (e.g., V), respectively (see Figure 3). The Case Filter is then applied by checking the co-occurrence of the attributes `ca, govern,` and `cm`.

3.3 Implementation of Parameters

While the principles described in the previous section are intended to be language independent, the structure of each grammar network in figure 1 is too language-specific to be applicable to languages other than the one for which it is designated. The most obvious language-specific feature is the ordering of head links with respect to complement links; in the graphical representation link ordering of this type is indicated by the starting points of links, e.g., C precedes IP under Cbar since the link leading to C is to the left of the link leading to IP. In the English network, all phrasal heads precede their complements. However, in head-final languages such as Korean, the reverse order is required. In order to capture this distinction, we incorporate the parameterization approach of (Dorr, 1993a) into the message-passing framework so that grammar networks can be automatically generated on a per-language basis.

The reason the message-passing paradigm is so well-suited to a parameterized model of language parsing is that, unlike head-driven models of parsing, the main message-passing operation is capable of combining two nodes (in any order) in the grammar network. The result is that a head-final language such as Korean is as efficiently parsed as a head-initial language such as English. What is most interesting about this approach is that model is consistent with experimental results (see, for example, (Suh, 1993)) which suggest that constituent structure is computed prior to the appearance of the head in Korean.

The remainder of this section describes our approach to parameterization of each subtheory of grammar described in the last section; we conclude with a summary of the syntactic parameter settings for English and Korean.

3.3.1 $\overline{\text{X}}$ Theory

$\overline{\text{X}}$ theory assumes that a constituent order parameter is used for specifying phrasal ordering on a per-language basis:

(2) **Constituent Order:** The relative order between the head and its complement can vary, depending on whether the language in question is (i) head-initial or (ii) head-final.

The structure above represents the relative order observed in Korean, i.e., the head-final parameter setting (ii). In English, the setting of this parameter is (i). This ordering information is encoded in the grammar network by virtue of the relative ordering of integer id's associated with network links. Other types of parameters encoded in the grammar network are those pertaining to basic categories (i.e., possible replacements for X in (1) above), pre-terminal categories (e.g., determiner), potential specifiers, and adjuncts for each basic category.

3.3.2 Trace Theory

In general, adjunct nodes are considered to be barriers to movement. However, Korean allows the head noun of a relative clause to be construed with the empty category across more than one intervening adjunct node (CP), as shown in the following:

(3) [$_{\text{CP}}$ [$_{\text{CP}}$ t_1 t_2 kyengyengha-ten] hoysa$_2$-ka manghayperi-n] Bill$_1$-un
yocum uykisochimhay issta
[$_{\text{CP}}$ [$_{\text{CP}}$ managed-Rel] company-Nom is bankrupt-Rel] Bill-Top
these days depressed is
'Bill, who is such a person that the company he was managing has been bankrupt, is depressed these days'

The subject NP 'Bill' is coindexed with the trace in the more deeply embedded relative clause. If we assume, following (Chomsky, 1986b), that relative clause formation involves movement from an inner clause into an outer subject position, then the grammaticality of the above example suggests that the Trace theory must be parameterized so that crossing more than one barrier is allowed in Korean. Our formulation of this parametric distinction is as follows:

(4) **Barriers:** (i) only one crossing permitted; (ii) more than one crossing permitted.

In English the setting be (i); in Korean the setting would be (ii).

3.3.3 Case Theory

In general, it is assumed that the relation between a case assigner and a case assignee is biunique. However, this assumption rules out so-called multiple subject constructions, which are commonly used in Korean:

(5) John-i phal-i pwureciessta
-Nom arm-Nom was broken
'John is in the situation that his arm has been broken'

The grammaticality of the above example suggests that Nominative Case in Korean must be assigned by something other than tensed Infl. Thus, we parameterize Case Assignment as follows:

(6) **Case Assignment:** Accusative case is assigned by transitive V; Nominative case is assigned by (i) tensed Infl; (ii) IP predication.

In a biunique case-assignment language such as English, the setting for Nominative case assignment would be (i); in Korean, the setting would be (i) and (ii).

3.3.4 Summary of Parameter Settings for English and Korean

We have just seen that certain types of syntactic parameterization may be captured in the grammar network (e.g., $\overline{\mathrm{X}}$ parameters such as constituent order). In addition to these, there are syntactic parameters (e.g., Trace and Case) that must be programmed into the message-passing mechanism itself, not just into the grammar network.

We adopt the syntactic parameters of (Dorr, 1993a) and focus on the automatic generation of the corresponding grammar networks from $\overline{\mathrm{X}}$ parameter settings (i.e., Basic Categories, Pre-terminals, Constituent Order, Specifiers, and Adjunction). The parameter compilation algorithm consists of two steps: the first defines the basic structural description (i.e., bar-level nodes) using the Basic Categories, Constituent Order, and Pre-terminals parameters; and the second defines specifier and adjunct links using the Specifiers and Adjunction parameters. The English and Korean grammar networks in figure 1 are the result of executing this algorithm on the Korean $\overline{\mathrm{X}}$ parameter settings.

3.4 Results of Time Test Comparisons

As a broad-coverage system, PRINCITRAN is very efficient. The parsing component (PRINCIPAR) processes real world sentences of length 20 to 30 words from sources such as Wall Street Journal within a couple of seconds. The complexity of the current version of the system has not yet been formally determined. However, we have demonstrated that the efficiency of the system is not purely a result of using an efficient programming language (C++); this has been achieved by running experiments that compare the performance of the parser with two alternative CFG parsers. Since PRINCIPAR has a much broader coverage than these alternative approaches, the absolute measurements do not provide a complete picture of how these three systems compare. However, the most interesting point is that the trends of the three performance levels relative to sentence length are essentially the same. If PRINCIPAR had an **average** case complexity that was exponential relative to sentence length, but had only managed to be efficient because of the implementation language, the sentence length vs. performance curve would clearly be different from the curves for CFG parser which are known to have a **worst** case complexity that is polynomial relative to sentence length.

The two CFG parsers used for comparison are: a C implementation of Tomita's parser by Mark Hopkins (University of Wisconsin-Milwaukee, 1993). The test sentences are from (Tomita, 1986). There are 40 of them. The sentence lengths vary from 1 to 34 words with an average of 15.18. Both CFG parsers use the Grammar III in (Tomita, 1986, p.172–6), which contains 220 rules, and small lexicon containing only the words that appear in the test sentences. The lexicon in PRINCIPAR, on the other hand, contains about 90,000 entries extracted from machine readable dictionaries.

Tomita's parser runs about 10 times faster than PRINCIPAR; Lin and Goebel runs about twice as fast.[6] To make the parsing time vs. sentence length distribution of these three parsers more comparable, we normalized the curves; the parsing time of each of the CFG parses was multiplied by a constant so that they would have the same average time as PRINCIPAR. The adjusted timings are plotted in Figure 4. These results show that PRINCIPAR compares quite well with both CFG parsers.

3.5 Implications for Machine Translation

Our ultimate objective is to incorporate the parameterized parser into an interlingual MT system. The current framework is well-suited to an interlingual design since the linking rules between the

[6] This is much slower than the result reported in (Lin and Goebel, 1993). This is because we are using a simulated version of the original system (i.e., PRINCIPAR without attribute values in the grammar and lexicon). The overhead for operations such as memory allocation/deallocation is higher in the simulated version, despite the fact that attribute values are not used.

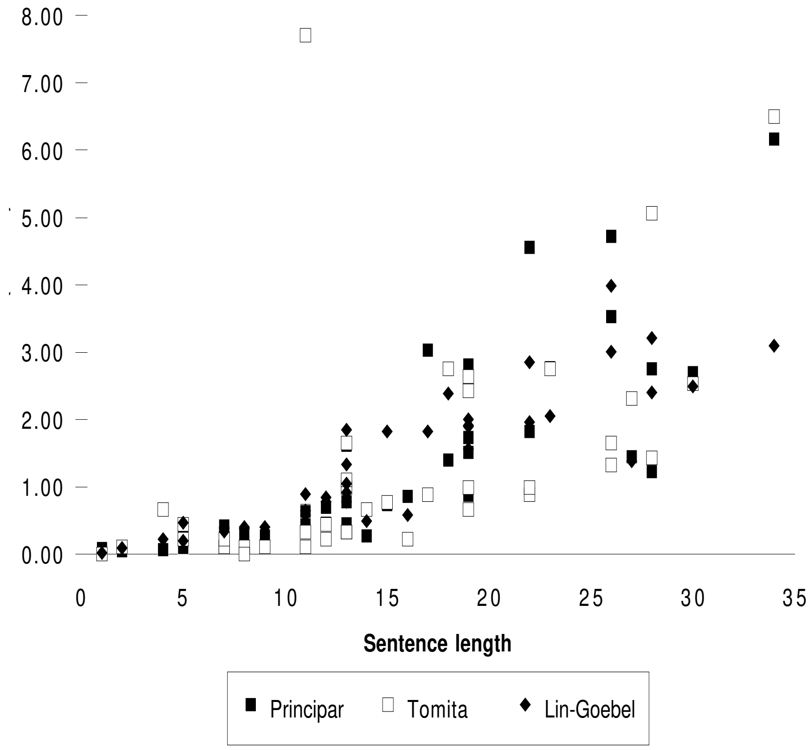

Figure 4: Adjusted Timings of Three Parsers

syntactic representations given above and the underlying lexical-semantic representation are well-defined (Dorr, 1993b). We adopt the Lexical Conceptual Structure (LCS) of Dorr's work and use a parameter-setting approach to handle well known, translationally divergent sentences (Dorr, 1990).

A parametric approach to mapping between the interlingua and the syntactic structure in English, Spanish, and German is described by (Dorr, 1990), (Dorr, 1993a). We present analogous examples here for English and Korean:[7]

(7) **Structural Divergence:**
E: John married Sally
K: John-i Sally-wa kyelhonhayssta
-Nom -with married
'John married with Sally'

(8) **Conflational Divergence:**
E: John helped Bill
K: John-i Bill-eykey towum-ul cwuessta
-Nom -Dative help-Acc gave
'John gave help to Bill'

(9) **Categorial Divergence:**
E: John is fond of music
K: John-un umak-ul coahanta

[7] The examples presented here are not necessarily geared toward demonstrating the full capability of the parser, which handles many types of syntactic phenomena including complex movement types. (See (Lin, 1993) for more details.) Rather, these examples are intended to illustrate that the parser is able to handle translationally contrastive sentences equally efficiently.

	Parse	Time
E:	[CP [Cbar [IP [NP [Nbar [N John]]] [Ibar [VP [Vbar [V_NP married] [NP [Nbar [N Sally]]]]]]]]	.15 sec.
K:	[CP [Cbar [IP [NP [Nbar [N John-i]]] [Ibar [VP [Vbar [PP [Pbar [NP [Nbar [N Sally]]] [P wa]]] [V_PP kyelhonhaysstа]]]]]]	.12 sec.
E:	[CP [Cbar [IP [NP [Nbar [N John]]] [Ibar [VP [Vbar [V_NP helped] [NP [Nbar [N Bill]]]]]]]]	.10 sec.
K:	[CP [Cbar [IP [NP [Nbar [N John-i]]] [Ibar [VP [Vbar [PP [Pbar [NP [Nbar [N Bill]]] [P eykey]]] [NP [Nbar [N towum-ul]]] [V_PP_NP cwuessta]]]]]]	.19 sec.
E:	[CP [Cbar [IP [NP [Nbar [N John-un]]] [Ibar [VP [Vbar [V_AP is] [AP [Abar [A fond] [PP [Pbar [P of] [NP [Nbar [N music]]]]]]]]]]]]	.12 sec.
K:	[CP [NP [0] [Nbar [N John-un]]] [Cbar [IP t [0] [Ibar [VP [Vbar [NP [Nbar [N umak-ul]]] [V_NP coahanta]]]]]]	.07 sec.

Figure 5: Parameterized Parsing of English and Korean Divergence Examples

-Nom music-Acc like
'It is John (who) likes music'

We ran the parameterized parser on both the English and Korean sentences shown here. The results shown in Figure 5 were obtained from running the program on a Sparcstation ELC. In general, the times demonstrate a speedup of 2 to 3 orders of magnitude over previous principle-based parsers on analogous examples such as those given in (Dorr, 1993b). Even more significant is the negligible difference in processing time between the two languages, despite radical differences in structure, particularly with respect to head-complement positioning. This is an improvement over previous parameterized approaches in which cross-linguistic divergences frequently induced timing discrepancies of 1-2 orders of magnitude due to the head-initial bias that underlies most parsing designs.

A preliminary investigation has indicated that the message-passing paradigm is useful for generation as well as parsing, thus providing a suitable framework for bidirectional translation. The algorithm for generation is similar to that of parsing in that both construct a syntactic parse tree over an unstructured or partially structured set of lexical items. The difference is characterized as follows: in parsing, the inputs are sequences of words and the output is a structure produced by combining two adjacent trees into a single tree at each processing step;[8] in generation, the inputs are a set of unordered words with dependency relationships derived from the interlingua (LCS). The generation algorithm must produce structures that satisfy the same set of principles and constraints as the parsing algorithm.

Three areas of future work are relevant to the current framework: (1) scaling up the Korean dictionary, which currently has only a handful of entries for testing purposes;[9] (2) the installation of a Kimmo-based processor for handling Korean morphology;[10] and (3) the incorporation of non-structural parameterization (i.e., parameters not pertaining to $\overline{\mathrm{X}}$ theory such as Barriers and Case

[8] The LCS composition routine described in (Dorr, 1992) derives the interlingua from the resulting syntactic representation.

[9] Our English dictionary has 90K entries, constructed automatically by applying a set of conversion routines to OALD entries. We have begun negotiations with the LDC for the acquisition of a Korean MRD for which we intend to construct similar routines.

[10] The English dictionary used by the message-passing system contains all morphological derivatives of every word. This approach would be impractical for Korean since the morphology is significantly richer.

Assignment).

3.6 LCS Composition

We are in the process of building the routines that convert the output of the parser described above into the LCS representation. The specification for the parser output is given in the form of a BNF in Appendix B. Sample syntactic and LCS dictionaries for Korean are given in Appendix C. Sample parse-tree and LCS outputs for a set of Korean and English sentences are given in Appendix D.

One of the advantages to the LCS representation is that it is easily extensible to a speaker-hearer model of discourse, as has been demonstrated by Dr. Martha Palmer and her NLP team at the University of Pennsylvania. A preliminary set of discourse structures for military messages is given below; these are samples that will ultimately be built automatically in the next phase of this project.

1. Message from "ops" to "bde"

   Text: COMMO IS UP.

   LCS:

```
[event CAUSE
        [thing SPEAKER]
        [state BE-ident
          [thing HEARER]
          [position AT-ident
            [property KNOW-IS-TRUE
              [state BE-ident
                [thing RADIO]
                [position AT-ident [property FUNCTIONAL]]
                [time NOW]
              ]
            ]
          ]
          [time [NOW to PLUS-INFINITY]]
        ]
        [time NOW]
      ]
```

2. Message from "ops" to "bde"

   Text: PLEASE SEND A CURRENT CMDRS REPORT ON ALL UNITS.

   LCS:

```
[event REQUEST
       [thing SPEAKER]
       [event CAUSE
         [thing HEARER]
         [event GO-poss
           [thing CURRENT-CMDRS-REPORT-ON-ALL-UNITS]
           [path TO-poss [position AT-poss [thing SPEAKER]]]
           [time [NOW to PLUS-INFINITY]]
         ]
         [time [NOW to PLUS-INFINITY]]
```

```
      ]
      [time NOW]
    ]
```

3. Message from “ops” to “bde”

   Text: DID YOU RECEIVE OUR LAST REQUEST?

   LCS:

```
[event REQUEST
       [thing SPEAKER]
       [event CAUSE
         [thing HEARER]
         [state BE-ident
           [thing SPEAKER]
           [position AT-ident
             [property KNOW-IS-TRUE
               [event GO-poss
                 [thing NOMINALIZATION
                   [event REQUEST'
                     [thing SPEAKER]
                     [event CAUSE' [thing X] [state/event Y]]
                     [time Z]
                   ]
                 ]
                 [path TO-poss [position AT-poss [thing HEARER]]]
                 [time [MINUS-INFINITY to NOW]]
               ]
             ]
           ]
           [time NOW]
         ]
         [time NOW]
       ]
       [time NOW]
     ]
```

4 Translation Issues

Upon completion of the previous annual report (January, 1994), we investigated certain issues that were left unresolved with respect to the more difficult translations between Korean and English. Some of these issues are addressed here.

4.1 Paraphrases and Resumptive Pronouns

In the previous report, some of the translations seemed to be paraphrases. An example of such a case is the following:

(10)
```
[cp[cp e1 e2  kyengyengha-ten]  hoysa2-ka  manghayperi-n]
                  managed-Rel  company-Nom is bankrupt-Rel
Bill1-un  yocum        uykisochimhay issta
```

```
     -Top these days  is depressed
  'Bill is such a person that the company which was managed
     by him has been bankrupt, and he is depressed these days'
```

Here, the gloss "Bill is such a person that the company which was managed by him has been bankrupt, and he is depressed these days" is the closest translation that we were able to provide. Note that one might mistakenly attempt to translate this as: "Bill, who managed the company which was bankrupt, is depressed these days" since this is a much more natural translation. This would be incorrect, however.

Consider these two cases more carefully:

```
(11) [[t2 t1 kyengyengha-ten] hoysa1-ka    mangha -n]     Bill2-un...
                managed-Rel  company-Nom was bankrupt-Rel Bill-Top
      (Literally): 'Bill2, (who2) the company1 (which1) HE2 was managing
         has been bankrupt, ...'
```

```
(12) [t2 [t1 mangha-n]   hoysa1-lul kyengyenghayss-ten] Bill2-un...
       was bankrupt-Rel company-Acc     managed -Rel   Bill-Top
      'Bill2, who2 managed the company1 which1 was bankrupt,...'
```

In both of these there are two relative clauses, one embedded inside the other. Also, in both cases, the more deeply embedded relative clause is modifying 'hoysa'(company) and the other one is modifying 'Bill'. However, there is a crucial difference between these two sentences. In the first one, it is Bill who is bankrupt. In the second one, it is the company that is bankrupt. Note that in the second case, 'Bill' can be associated with its trace (=t2) across only one S boundary. However, 'Bill' in (17) cannot be associated with its trace in that way. Rather, it is associated with its trace across two S boundaries.

In this case, it seems there is no direct English translation: associating a trace with its antecedent across more than one bounding node results in ungrammaticality in English (given that the trace results from a single application of movement). In a later version of this, we changed the translation to be less awkward, using the pronoun "whose" and a resumptive pronoun (as suggested by Yaeger):

```
Bill, whose company he was managing is bankrupt [so he] is depressed
these days.
```

However, given the device we are using for parsing (described in the last section), it seems there is no computational inexpensive way of restructuring the phrase so that it uses the resumptive pronoun. Such a modification would require the repositioning of traces, which is heavily computation-intensive.

4.2 Use of Chinese Characters for Semantic Distinctions

In the previous report, the following example was presented as a divergence:

```
(13) K: John-un       khi-ka     khuta
            -Top  height-Nom    is big/huge
        'John is tall'
     E: John is tall
```

In translating, there are important semantic distinctions between Korean words and their Chinese character counterparts. In the above example, the Korean word "khuta" is a predicate adjective that stands on its own. It would not be possible to use the Chinese analogue, "tay" because the latter is exclusively used as a modifier of a noun.

Often it is difficult to distinguish between the two cases:

(14) tay-cicin / cicin
     big-earthquake earthquake

(15) tay-hak / *hak
     big-study(?) study

In the first example above, 'tay' is a modifier of 'cicin'(earthquake). Crucially, 'cicin' can be used without 'tay' (as an independent word). This is not the case in the second example. The word 'tay-hak', which means 'university', cannot be separated as 'tay'+'hak'. (In other words, 'hak' alone cannot be used for representing 'study'.) In this case, 'tay' seems to be a word in its own right, i.e., an idiomatic meaning results. We may simply distinguish between such by considering 'tay' to be a modifier in the former but not in the latter.

On the other hand, the Korean word 'khu-'(big) can be used not only as a modifier but also as a predicate. In the former usage, it combines with adnominalizer '-n', resulting in 'khu-n'. In the latter usage, it combines with declarative marker'-ta' (or interrogative marker '-ni'), resulting in 'khu-ta' (or 'khu-ni').

4.3 Syntactic Structure of Korean

During the course of our study of the structure of Korean, we have attempted to address the issue of the necessity of a CP (clausal) node. The issue of whether Korean has a CP has been very controversial. We claim that a CP is, indeed, necessary and that there exist several empirical and conceptual arguments to support this.

The first argument is related to the notion of barrierhood. In (Chomsky, 1986a), the CP node can be a barrier, but IP cannot be a barrier by itself. IP can be a barrier only by inheritance from the maximal projection(=XP) it is dominating. If we hypothesize that there is no CP in Korean, IP must be a barrier, given that (at least) some movement out of certain clause creates ungrammaticality. This in turn leads to the (unmotivated) asymmetry between English and Korean: English IP is not a barrier by itself while Korean IP is. But if we assume that there is CP in Korean, then we need not resort to such an unmotivated asymmetry.

In Korean extraction phenomena in certain constructions (like the topic construction) are not accounted for by Subjacency. However, we still seem to need the notion of barrierhood. So, we adopt the assumption that there is a CP node in Korean and that non-L-marked CP is a barrier.

Another argument for the necessity of CP comes from the structure of relative clause or adverbial clause.

```
(16) [ t1   Mary-lul coaha-nun] salam1
            Mary-Acc  like-Rel   man
       'the man who likes Mary'
```

If there is no CP in Korean, then there is no site for accommodating the relative clause marker '-nun'; it cannot be placed in the Spec of IP or IP adjoined position. Since a relative clause should be distinguished from a complement clause, for instance, we need a place for accommodating a category which signals the type of clause, i.e., complementizer, relative clause marker, adverbial clause marker etc. This means that we must have CP (above IP).

4.4 Scrambling as Case-Driven Obligatory Movement

Another area that we have investigated intensively since the previous report is the issue of scrambling. In (Lee, 1993), movement is not optional, but is a consequence of case-driven obligatory movement. Some points about her proposal are the following:

(A) Scrambling is adjunction.

(B) Scrambling is A-movement: It exhibits properties of A-movement (in terms of Binding and Weak Crossover phenomena).

(C) Consequently, adjoined positions are A-positions in Korean.

(D) Scrambling is a consequence of case-driven movement. (Assuming the VP-internal subject hypothesis, all the arguments move out of VP to be assigned case. As long as case licensing conditions are met, arguments may be arranged in any order.)

There are some consequences and implications of the above proposal. The first is that Theta-role assignment is completely dissociated from case assignment. The second is that even a canonical word order sentence involves obligatory movement (scrambling). Finally, this proposal predicts a a parametric difference between English and Korean. The difference is the level at which accusative case is assigned: In Korean, accusative case is licensed at S-structure, while in English it is done at LF. (Such a difference is reduced to the level at which verb raising to INFL takes place.)

We adopt this proposal and take it one step further by proposing an algorithm that will allow us to bind empty arguments corresponding to scrambled elements. In particular, if an argument is scrambled out of a clause, we can re-claim the argument inside of clause by using a trace or a crossing branch:

(17)
```
John said that Mary likes Tom.
John-top Mary-nom Tom-acc like+comp said. (basic order)
Tom-acc John-top Mary-nom like+comp said. (scrambled, Tom is outside)
```

If it's coindexed, we can recover the argument by having a trace or by allowing more than one terminal point to the same NP:

(18)
```
John said that () likes Tom.
John-top Mary-nom like+comp said.

Analysis 1:
John-top1 e1 Mary-nom like+comp said.

Analysis 2:
e1 John-top1 Mary-nom like+comp said.
```

If there is no preceding NP with the case required by a particular verb, then topic-marked NP will be converted to an accusative-marked NP:

(19)
```
John said that Mary likes ().
John-top Mary-nom (here need NP-acc) like+comp said.

Analysis: John-top1 Mary-nom e1 like+comp said.
```

Note: John-top1 is co-indexed with e1 which need an accusative NP antecedent.

We are in the initial stages of designing an algorithm that takes the above points into consideration. This will be installed in the parser during the next phase of the project.

4.5 Empty Categories

As part of our work to reduce the complexity of movement and scrambling, we have had to face the difficulty of recoverability of Empty Categories (EC's). On approach that has been suggested is that of (Egedi et al., 1994), in which the general rule that is used is: 'Choose the closest topic that matches the semantic constraints on the elided arguments'. This approach is computationally efficient in that the antecedent of EC is recovered without delaying the actual computation of sentence structure. However, there is a problem with rule. In particular, there are frequently cases where the closest topic cannot be the antecedent of the EC, as seen in the following:

```
(20) Mary-ka   [John-i  e  hyeppakhayssta-ko]  malhayssta
     Mary-Nom  John-Nom   threatened-Comp       said
     'Mary said that John threatened her.'
```

The algorithm is to search through the list for each dropped argument, matching the semantic feature constraints from the verb against the features on the NP until the appropriate NP is found. Roughly, the structure assumed for the above sentence (with English words inserted) would be:

```
            S
          /   \
     Mary_1     VP
               /   \
             S    said
           /  \
     John_2    VP
              / \
             e   threatened
```

In this example, the empty object in the embedded clause must be coreferential with the matrix subject (Mary), although the closer NP to the empty object is the embedded subject (John). (cf: In their paper, only topic NPs have been discussed. However, their general rule can be (and must be) applicable to NPs other than topic-marked ones. In (1), there is no topic-marked NP; thus, the antecedent of EC should be sought among non-topic NPs.) In fact, as discussed in (Suh, 1994), empty objects in Korean tend to be bound by an NP outside of the current clause probably because they are pronominals constrained by the Binding Condition B. Such a grammatical principle should be implemented in in order to make correct predictions.

Note that one could argue that there is a semantic property associated with "threaten" (i.e., that one cannot threaten oneself) that forces the correct binding to fall out. However, it is easy to come up with a counterexample. Consider the following structure, which is legal in Korean:

```
            S
          / \
     Mary_1 VP
             / \
            S said
          / \
     John_2 VP
             / \
            e liked
```

The algorithm above would obtain an incorrect result, i.e., the empty category, e, would be incorrectly bound to the closest NP John, yet it should be bound to the matrix subject NP Mary. Note that it is possible for someone to "like himself". In particular, it is not the case that what is responsible for the binding of EC from outside of the current clause is a semantic factor.

Consider another example in which it is semantically possible for the EC to be bound by the closest NP, i.e., the EC still should be bound from outside of the current clause:

```
(21) John-i [Mary-ka e piphanhayssta-ko/ chingchanhayssta-ko]  malhayssta
     John-Nom Mary-Nom EC criticized-Comp/ praised-Comp         said
     'John said that Mary criticized/ praised  him (or someone else).'
```

It is semantically possible for someone to criticize or praise himself/herself. But the example above does not allow such a semantic interpretation, and the EC should find its antecedent outside of the current clause. Consequently, 'John' (or somebody else available from the context) is coindexed with the EC. It seems clear from this example that such an EC is a pronominal subject to Binding Condition B (or a variable subject to Binding Condition C).

5 Temporal/Aspectual Analysis

In the next phase of the project, we will be collaborating with Dr. Mari Olsen from Northwestern University in order to install temporal/aspectual knowledge into our semantic representations. As a preliminary study, in preparation for this upcoming collaboration, we have worked on a definition of basic tense structures for Korean which we expect to tie into the work of (Olsen, 1994). In addition, we have undertaken the task of investigating a number of aspectual paradigms and have subsequently attempted to distinguish among different levels of representation that are affected by aspectual information. This "layered" approach is one that will mesh well with the work of Dr. Olsen, in which temporal, lexical, and grammatical aspect are carefully separated into different tiers.

5.1 Possible Tenses in English and Korean

According to (Hornstein, 1990), in all of the languages of the world, there are (logically) 24 possible basic tense structures (BTS):

```
Present: S,R,E  S,E,R  R,S,E  R,E,S  E,S,R  E,R,S
Past:    E,R_S  R,E_S
Future:  S_R,E  S_E,R
Present perfect: E_S,R  E_R,S
Past perfect:    E_R_S
Future perfect:  S_E_R  S,E_R  E_S_R  E,S_R
Distant future:  S_R_E
Future in past:  R_S,E  R_E,S  R_S_E  R_E_S
Proximate future:S,R_E  R,S_E
```

BTSs are composed of an SR relation and an RE relation: SR and RE are either associated or separated, and are strongly ordered linearly. However, S and E are not related to each other in these ways except derivatively. Given this, we can cut down the number of possible tenses to 16:

```
S,E,R = *E,S,R
R,E,S = *R,S,E
S_E_R = *S,E_R = *E_S_R = *E,S_R
R_S,E = *R_E,S = *R_S_E
S,E_R = *E,S_R
```

Using the framework above, we developed a scheme in which English and Korean morphemes were associated with BTS representations. We will look at each of these languages, in turn:

(A) English:

(a) i. present morpheme: associate S and R: `S,R`
   ii. past morpheme: R removed to left of S: `R_S`
   iii. future morpheme: R removed to right of S: `S_R`

(b) i. +have: E removed to left of R: `E_R`
   ii. -have: E and R associated: `E,R` or `R,E`

**There is no specific morpheme that signals the association of R and E.

(B) Korean:

(a) i. present morpheme: associate S and R: `S,R`
   ii. past morpheme: R removed to left of S: `R_S`
   iii. future morpheme: R removed to right of S: `S_R`

(b) E and R are generally associated(=E,R or R,E)

**In a given BTS, if linear order is not intrinsically determined, assume that the linear order of RE is identical to the linear order of SR. Also, morphemes unambiguously determine unique mappings.

In the final inventory of 8 tenses (shown here), Korean only has the first three. This is because E and R are generally associated (unless specified otherwise).

```
Present: (S,R)o(R,E) = S,R,E  (i)
         (R,S)o(E,R) = E,R,S  (ii)
Past:    (R_S)o(E,R) = E,R_S
Future:  (S_R)o,(R,E) =  S_R,E
Present perfect: (S,R)o(E_R) = E_S,R  (i)
                 (R,S)o(E_R) = E_R,S  (ii)
Past perfect:    (R_S)o(E_R) = E_R_S
Future perfect:  (S_R)o(E_R)
Future in past:  (R_S)o(R_E)
Proximate future:(S,R)o(R_E) = S,R_E  (i)
                 (R,S)o(R_E) = R,S_E  (ii)
```

We now consider some examples. The first case is Past vs. Past Perfect. The latter is is obtained by "reduplication" of the past tense marker:

```
(22) John-un Mary-lul  piphanhay-ss-ta
         -Top    -Acc  criticize-Pst-Ind
     'John criticized Mary.'
          E,R_S
```

```
(23) John-un Mary-lul piphanhay-ss-ess-ta
         -Top    -Acc criticize-Pst-Pst-Ind
     'John had criticized Mary.'
     E_R_S
```

The second case is Future vs. Future Perfect. The latter interpretation is available when the topic marker is attached to the temporal adverbial phrase, which amounts to the R point:

```
(24) John-un  han     si-ey  swukcey-lul kkutnay-l kes-ita
         -Top one o'clock-at homework-Acc   finish-will
     'John will finish his homework at 1 p.m.'
     S_R,E
        |
        1 p.m.
```

```
(25) John-un  han     si-ey-nun  swukcey-lul kkutnay-l kes-ita
         -Top one o'clock-at-Top homework-Acc  finish-will
     'John will finish his homework by 1 o'clock.'
     'John will have finished his homework at 1 o'clock.'
     S_E_R
         |
       1 p.m.
```

Finally, consider the case of Present perfect vs. Past. Present perfect meaning is represented through the use of particular expressions involving (auxiliary) verbs.

```
(26) John-un New York-ulo  ka pely-ess-ta/    ka-ko    ep-ta
         -Top         -to   go  Aux-Pst-Ind   go-conj  not exist-Ind
     'John has gone to New York.'
     E_S,R
```

```
(27) John-un New York-ulo  ka-ss-ta
         -Top         -to   go-Pst-Ind
     'John went to New York.'
     E,R_S
```

Due to the lack of a grammatical device for representing perfect tense (corresponding to English auxiliary 'have'), emphasizing the R point seems necessary in Korean in order to represent perfective meaning explicitly. Case (27) is such an instance, where the topic marker is attached (for making the R point salient) to the PP referring to the R point. (Such an usage of topic marker is also observed in past perfect.)

5.2 The Place of Aspect in an IL-Based MT System

In an effort to place our work in the wide range of analyses available in the literature we have adopted the following definitions of what "aspect" is:

(a) semantic vs. grammatical aspect:
"aspect" may appear in sentences lexically and morphologically as overt grammatical evidence for the underlying or semantic notion.

(b) situation vs. perspective aspect:
"aspect" may refer to the internal structure of a situation (states, processes, events/transitions) or to how much of the full situation is being focused on (perfective, imperfective)

(c) temporal vs. agentive/volitional aspect: "aspect" has been taken to be a temporal property by some whereas others have defined it more broadly to include agentive or volitional properties

With respect (b), we represent situation aspect using a temporal property definition of aspect. We are focusing specifically on those properties that are derived from spatial information. (For example, "John ran in the park" conveys a Process, and so does "John ran through the park". However, "John ran to the park" does not, it conveys a bounded Event.) This gives us the opportunity to test the robustness of our spatial representations and to start building a micro-theory of aspect for translation.

With respect to (a), any MT system must be able to capture both semantic and grammatical aspect. The mapping between semantic and surface or grammatical aspect is language-specific, whereas the semantic aspect itself in an IL-based MT system must be language-independent. Clearly the development of the mapping and the representation itself are related. The representation for aspect may either be the primary focus with the mapping developed secondarily, or the mapping may be primary, in which case it may drive the representation of aspect. We have chosen the first of these two approaches.

At the IL level of representation of a sentence we are translating, the MT system has access to two classes of information:

(1) language-specific information used to map from the source language lexicon entries and parse tree to the IL, and from the IL to target language lexicon entries and generation tree

(2) language-independent, IL-theory-internal definitions to map IL primitives into some subset of primitives in knowledge representation (KR) system

Though strictly speaking we take "aspect" to be a property of IL structures of type SITUATION (as opposed to type THING, SPACE, etc.), the information that goes into deriving situation-aspect, call it "i-aspect", may appear at any level in the MT system.

Any MT system that claims to be "interlingua-based" must of course justify which levels exist in the design and what constitutes evidence for the representational information at each level. As evident above, we have an IL and a KR level of representation. It is our working assumption that i-aspectual information corresponds (possibly through mappings between levels) to primitives at the KR level.

We can identify 3 ways in which i-aspectual information arrives in an IL structure. First, if i-aspectual information originates (i.e. is available) lexically or morphologically, then the MT system has access to it by virtue of it being "coded" in the system SL or TL lexicon. For example, the word "repeatedly" requires including i-aspectual information in its lexicon entry.

Second, i-aspectual information may be part of the definition of an IL primitive. The information does not appear in the SL or TL lexicons directly, only indirectly via the primitive. For example, the word "through", as in "John ran through the park", has the IL primitive VIA in its lexicon entry. It is only in the IL-theory-internal definition of VIA that we have i-aspectual information in terms of KR primitives for VIA, and hence for "through". So the MT system has access to it by virtue of it being "coded" in the IL-theory-internal lexicon, where IL primitives are listed as entries together with their KR definition.

Third, i-aspectual information may need to be derived because it is not available in either of the above ways even though the sentence being translated is considered grammatical. The derivation is contingent on the existence of "coercion" or "construal" rules that are invoked to complete an IL structure. For example, consider the sentence "she lives through the tunnel". This means she lives at the other end of the tunnel. The issue here is where does this particular interpretation of the sentence come from given that the sentence conveys a STATE situation even though the prepositional phrase contains PATH information, rather than the needed PLACE information. We could fix this lexically and define one meaning of "through" to be an augmentation of its PATH meaning (with its VIA primitive, as above) into a PLACE meaning, with an IL primitive such as AT-END-OF. (One limitation to this approach is that it may then overgenerate interpretations of "through" in cases where this meaning is not applicable.) Alternately, we could craft a construal rule that allows for such an augmentation more generally when a STATE verb requires a PLACE argument that converts a PATH into a PLACE. This is a context-dependent construal that may be more be finely-tuned than the lexical approach.

5.3 Lexical and Grammatical Aspect: Olsen

Our work with Dr. Mari Olsen has allowed us to fine-tune our specification of the levels of aspectual knowledge described above. Over the next phase of the project, we will be applying Olsen's model to MT for incorporation of aspectual knowledge into the interlingua. Olsen's privative analysis of aspect and tense allows temporal semantics to be built up monotonically, from features associated

Automation Spectrum	Lexicon Tool
Manual	LCS Editor
Semi-Automatic	TEAM (Grosz et al., 1987) IRUS (Bates and Bobrow, 1983) TELI (Ballard and Stumberger, 1986) LUKE (Knight, 1991) User's LCS Editor Others: (Ginsparg, 1983), (Guida and Tasso, 1983), (Grishman, 1986), (Thompson and Thompson, 1983), (Templeton and Burger, 1983), etc.
Automatic	LEXICALL

Figure 6: Spectrum of Possibilities: Lexicon Building for NLP Applications

with the verb, to those marked by other lexical and grammatical constituents and the pragmatic context (see (Olsen, 1994)).

In conjunction with this work Dr. Olsen will be collaborating with the University of Maryland on how lexical aspect may be used to improve lexical choice in machine translation. For example, Spanish has many motion verbs in which the path of the motion event is encoded (e.g. entrar 'enter'); English, according to (Talmy, 1978), (Talmy, 1991), more frequently uses verbs 'lexicalizing', i.e. expressing, the manner of the motion rather than the path (e.g. run, roll). These generalizations may translate into a difference in the frequency of telic verbs: Spanish may have more [+telic] verbs (types or tokens) than English. In some cases, therefore Spanish to English translation may be described as a mapping of a Spanish [+telic] verb (such as entrar) onto an English atelic verb with a [+telic] complement (such as go in). However complexities in this mapping need to be explored before an algorithmic description is possible.

6 Lexicon Development and Semantic Classification of Verbs

In this section, we discuss an area related to the problem of lexical acquisition, i.e., semantic classification of verbs and implementation of tools for the development of lexicons based on this classification.

6.1 Semi-Automatic Tools for Lexicon Development

In the previous annual report, we described our LCS Editor, which provides an interface that allows LCS's to be built and modified. During editing, the current state of the LCS is continuously updated and displayed in two windows.

It is clear that, while the LCS Editor is useful, it still does not eliminate many of the aspects that lead to tedium in the construction of a dictionary. In this section, we discuss a progression of dictionary-construction ideas ranging from the LCS Editor that we've already built to a more advanced lexicon development tool. Figure 6 illustrates this range.

At one end of the lexicon-tool spectrum (Manual), the level of knowledge assumed by the dictionary builder must be relatively extensive, the level of tedium is high, and development time is lengthy; however, linguistic coverage is easily controlled by the builder (i.e., full coverage within a domain is well within the realm of possibility). At the other end of the spectrum (Automatic), there need not be any linguistic knowledge and development is automatic (i.e., not tedious or lengthy); however, there lexicon coverage is likely to have "holes," i.e., a lack of application-specific information due to the fact that the content of machine-readable resources cannot always be predicted in advance.

The use of the LCS Editor requires a great deal of linguistic knowledge because LCS's are more sophisticated than other types of knowledge that are generally encoded in NLP lexicons.

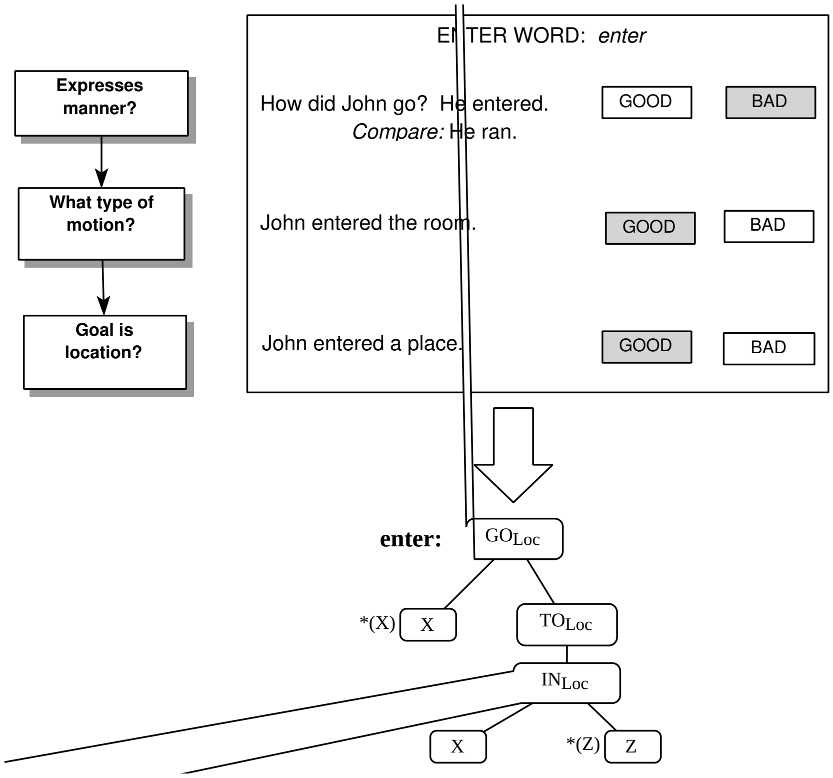

Figure 7: User's LCS Editor

However, the editor does provide certain types of automatic constraint satisfaction that fills in certain types of nodes while the user is building the representation. This is where the benefit of the LCS representation is most apparent: the LCS conforms to wellformedness conditions that can be applied automatically during the construction of a lexical entry.

Many lexicon development tools fall under the category of "Semi-Automatic," i.e., they have been built for users with no specialized training in linguistics. Most tools in this category still look much like the first such interface created for the TEAM project (Grosz et al., 1987). One of the most sophisticated tools in the Semi-Automatic category is the LUKE system by (Knight, 1991). This system comprises the lexical acquisition component of large NLP system called KBNL (Barnett et al., 1990). LUKE is connected to a knowledge-representation system called CYC (Lenat and Guha, 1990), a large-scale, common-sense, knowledge-based (KB) project. The drawback of this system is that the person who does editing must be KB expert and, in particular, must know about concepts available in the CYC system.

The approach used in systems like LUKE is clearly a step in the right direction toward automation of the lexicon construction process. While many of the ideas behind these systems have been borrowed in our approach to lexicon construction (e.g., the mapping from syntactic constituents to semantic roles), we have found that it is possible to avoid certain drawbacks such as the requirement that the lexicon developer be a KB expert.

The latest version of our LCS Editor is one that allows a non-linguist i.e., a "User" of the Foreign Language Tutoring system — generally assumed to be a foreign language instructor — to enter words into the lexicon without being forced to manipulate the LCS representation directly. (See figure 7.) This tool does not require as much linguistic knowledge as our already existing Lexicon Tool, but is still in the semi-automatic category as is the LUKE system. This approach has the advantage that it does not require a large knowledge base (or a KB expert). However, as in LUKE, the system presents easy-to-read sentences and the user is then asked to assign grammaticality judgements to these sentences. Note that the menu from the previous figure is given in parallel here to illustrate the correspondence between menu-driven questions and the relevant sentences. Once the user makes a judgement about these sentences, the information is then used to infer

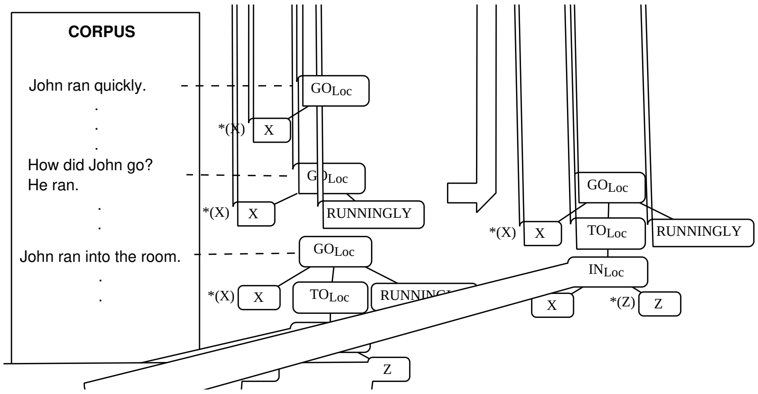

Figure 8: Toward Automatic Lexical Acquisition: Searching a corpus for sentences

information about the LCS representation.

The key to the success of this design is the use of an appropriate set of sentences that allow automatic inference information about the LCS. For example, the verb *enter* is different from *run* in that the former does not express manner of motion. This subtle distinction is brought out by asking the user to specify whether *How did John go?/He X-ed* is an appropriate sentence pair. For *run* this pair is appropriate since the "how" portion of the verb is inherent in its meaning — he went "runningly"; by contrast, for *enter* this pair is not appropriate since there is no "how" component of meaning.

If we push the above design one step closer to automated lexical acquisition, we might consider building a sentence-finding program that searches a textual corpus (i.e., a large body of sentences in a particular language) for the occurrence of sentences (or sentence pairs) that correspond to the ones in the presentation window of figure 7. Thus, instead of developing the appropriate set of sentences to isolate meaning, the task is now to devise a program that is able to select appropriate sentences from corpora. Figure 8 illustrates this modified design. In this example, the system is attempting to determine the LCS representation for the verb *run*. Note that each sentence that is found to be relevant in the corpus is analyzed for information that would add one component of meaning to the LCS representation. For example, the verb phrase *run quickly* in the first sentence indicates that this is a GO action; the addition of the manner component *runningly* can then be derived from the sentence pair, i.e., from the occurrence of a sentence that answers the question "how".

The key to the success of this type of automatic design is to use an extensive corpus (for each of the languages of interest) and a "smart" analyzer. Unfortunately, as shown by (Hogan and Levin, 1994), it is a non-trivial task to acquire even simple components of meaning, such as argument structure information, by means of syntactic extraction out of a corpus. Moreover, a corpus has inherent limitations; we have no guarantee that we will be able to find the precise sentences we need in order to predict LCS components of meaning. For example, even with a corpus as enormous as the Lancaster-Oslo-Bergen (LOB) corpus — 60,000 sentences — there is no occurrence of the construction "How did `X Verb`", with or without a subsequent response of the form "`X Verb`ed".

The difficulty of searching for *particular* constructions has led us to investigate a higher level of abstraction that utilizes the verb classification presented in (Levin, 1993). Our goal is to test Levin's hypothesis that the syntactic behavior of a word is fully semantically determined and, moreover, that this property holds across all languages. For example, the verb *run* is a manner-of-motion verb which, according to Levin, participates in the following syntactic alternations:

Path	The horse ran through/into/out of the stream
There Insertion	A horse ran out of the barn There ran out of the barn a horse
Locative Inversion	A horse ran out of the barn Out of the barn ran a horse
Measure Phrase	We ran 5 miles
Resultative Phrase	We ran ourselves into a state of exhaustion

Searching for alternations in a corpus has led to more promising results than searching for a small set of language-specific constructions. Part of the investigation of this approach has involved the compilation of a large matrix that crosses syntactic alternations (approximately 80) with semantic classes (approximately 192). Preliminary results (Corbin et al., 1994) have demonstrated that this approach is easy to implement; however, the mapping from semantic classes to LCS representations is non-trivial. Moreover, many of the syntactic alternations are specific to English, which means that we would have to build a table like this for every language for which we need a dictionary. Given that the work of (Levin, 1993) was the result of compilation of at least a decade of research articles and notes, it is clear that the construction of a table for every language is a non-trivial task.

Another concern is the danger of committing to specific alternations that might have radically different analogues in other languages.[11] (Levin, 1993) argues that when the same alternations *do* exist across languages, the same semantic classification is supported; however, (Mitamura, 1990) opposes this point of view, arguing that Japanese associates syntactic patterns with semantic classes, but that these semantic classes are not the same as those developed by Levin for English. She claims that syntactic alternations in Japanese are based primarily on case distinctions, for example, verbs associated with the *ga* and *kara* case markers are in the semantic class of *giving*. Since there is no semantic class of 'giving' of this type in (Levin, 1993), Mitamura argues that Levin's semantic classification is not correct.

In our own study of Korean (which is similar to Japanese in structure), we too have discovered that the syntactic tests for English do not apply directly to Korean verb classes. An example is the transitivity alternation, which is an important criterion for Levin's verb classification. This alternation is irrelevant to Korean because it is not possible for a Korean verb (such as *break*) to function as both a transitive and an intransitive verb. However, it is not clear that the tests that Mitamura was relying on are as systematic as they could be. In particular, Mitamura's grammatical assumption is that the notions of SUBJECT and OBJECT are primitive (as in Lexical Functional grammar). She provides the following examples of Alternations in English and Japanese:

```
(28) a. Brian supplied a blanket to John.
     b. Brian supplied John with a blanket.

(29) a. Brian-ga  John-ni    mouhu-o    sikyuusita
            -Nom     -Dat  blanket-Acc   supplied
        'Brian supplied a blanket to John.'
     b.*Brian-ga  John-o     mouhu-de    sikyuusita
            -Nom     -Acc  blanket-with   supplied

(30) a. kodomo-ga   kawa-de   oyoida
         child-Nom  river-in   swam
       'A/The child swam in the river.'
     b. kodomo-ga   kawa-o    oyoida
```

[11] Surprisingly, the analogues of many English syntactic alternations do actually exist in other languages. For example, the Australian language Warlpiri has a middle construction (i.e., *The bread cuts easily*).

```
      child-Nom  river-Acc   swam
      'A/The child swam in the river.'
```

Mitamura argues that the alternation above occurs with the movement verbs which take a locative phrase. (The marker -de indicates a general location where the movement is held, and the marker -o indicates the area covered by the movement.)

Our view is that the notions of 'subject'/'object' as primitives is highly suspect. The validity of the 'subjecthood'/'objecthood' test employed by Mitamura is questionable. Consider the following examples:

```
(31) a. John-i  Bill-ul   ttayli-ess-ta
            -Nom     -Acc     hit-Pst-Ind
        'John hit Bill.'

     b. *Bill-i  John-eyuyhay  ttaylie-ci-ess-ta
             -Nom     -by          hit-Pass-Pst-Ind
        'Bill was hit by John.'

(32) a. John-i  Bill-ul   kwuthaha-ess-ta
            -Nom     -Acc      hit-Pst-Ind
        'John hit Bill (repeatedly).'

     b. Bill-i  John-eyuyhay  kwutha-tangha-ess-ta
            -Nom     -by          hit -Pass-Pst-Ind
        'Bill was hit by John (repeatedly).'
```

Verbs such as ttaylita(hit), chita(strike), cwukita(kill), sata(buy), yokhata(curse), mannata(meet), topta(help), chingchanhata(praise) etc. take an accusative NP as their complement, which is NOT passivizable. According to Mitamura, such an accusative NP is not an object but an oblique NP, and the verbs taking such an NP as a complement would be distinguished from the verbs (with similar meaning) which take a passivizable complement. This is not the desired result. Thus, we dismiss the notions of subject, object, and oblique (as grammatical function), and concentrate on 'typical' case marking and 'typical' theta role of NPs in classifying verbs.

Levin's work does not assume that a verb in one language necessarily behaves the same as all of its equivalent forms in another language, especially if that verb has multiple verb senses. For example, there is no single Korean counterpart for all senses of the English verb *break*, so it should not be the case that *all* syntactic behaviors associated with the word *break* carry over to *all* possible counterparts in another language. We take the view adopted by Levin that, while a verb in one language might not have an exact counterpart in another language, the "basic meaning components" that define the lexical items of each language are the same.

Our approach to automatic lexicon construction relies heavily on this notion of "basic meaning components." If we can decompose Levin's semantic classification into primitive units of meaning, then we should be able to conduct a more thorough investigation of whether the proposed semantic classification is indeed applicable to other languages.

6.2 Automated Lexicon Construction: LEXICALL

We have designed a lexical acquisition program that relies on a systematic relation between syntactic behavior and "basic meaning components" for the construction of LCS representations. We take as our starting point the 3100 verbs that have already been classified in (Levin, 1993). Through the statistical process described below in section 7, we classified an additional 3100 verbs. Once we arrived at the full Levin-based classification for 6200 verbs, our acquisition algorithm used a