

Telicity as a Cue to Temporal and Discourse Structure in Chinese-English Machine Translation*

Mari Olsen

Microsoft
molsen@microsoft.com

David Traum

U Maryland
traum@cs.umd.edu

Carol Van Ess-Dykema

U.S. Department of Defense
carol@umiacs.umd.edu

Amy Weinberg

U Maryland
weinberg@umiacs.umd.edu

Ron Dolan

Library of Congress
rdolan@cfar.umd.edu

Abstract

Machine translation between any two languages requires the generation of information that is implicit in the source language. In translating from Chinese to English, tense and other temporal information must be inferred from other grammatical and lexical cues. Moreover, Chinese multiple-clause sentences may contain inter-clausal relations (temporal or otherwise) that must be explicit in English (e.g., by means of a discourse marker). Perfective and imperfective grammatical aspect markers can provide cues to temporal structure, but such information is not present in every sentence. We report on a project to use the lexical aspect features of (a)telicity reflected in the Lexical Conceptual Structure of the input text to suggest tense and discourse structure in the English translation of a Chinese newspaper corpus.

1 Introduction

It is commonly held that an appropriate interlingua must allow for the expression of argument relations in many languages. This paper advances the state of the art of designing an interlingua by showing how aspectual distinctions (telic versus atelic) can be derived from verb classifications primarily influenced by considerations of argument structure, and how these aspectual distinctions can be used to fill lexical gaps in the source language that cannot be left unspecified in the target language. Machine translation between any two languages often requires the generation of information that is implicit in the source language. In translating from Chinese to English, tense and other temporal information must be inferred from other grammatical and lexical cues. For example, Chinese verbs do not necessarily specify whether the event described is prior or cotemporaneous with the moment of speaking. While grammatical aspect information can be loosely associated with time, with imperfective aspect (Chinese 在 *zai*- and 着 *-zhe*) representing present time and perfective (Chinese 了 *le*) representing past time, (Chu,

1998; Li and Thompson, 1981), verbs in the past do not need to have any aspect marking distinguishing them from present tense verbs. This is unlike English, which much more rigidly distinguishes past from present tense through use of suffixes. Thus, in order to generate an appropriate English sentence from its Chinese counterpart, we need to fill in a potentially unexpressed tense.

Moreover, Chinese multiple-clause sentences may contain implicit relations between clauses (temporal or otherwise) that must be made explicit in English. These multiple-clause sentences are often most naturally translated into English including an overt expression of their relation, e.g., the "and" linking the two clauses in (1), or as multiple sentences, as in (2)).

- (1) 1965 年 前 , 我 国 总 共
1965 year before , our.country altogether
只 有 30 万 吨 的
only have 30 ten-thousand ton de
造 船 能 力 , 年 产 量 是 8
shipbuilding capacity , year output is 8
万 吨
ten-thousand ton

Before 1965 China had a total of only 300,000 tons of shipbuilding capacity and the annual output was 80,000 tons.

- (2) 这 8 万 吨 竟 包 括 了
this 8 ten-thousand ton actually include asp
517 艘 , 船 的 吨 位 是 很 低 的
517 cl , ship de tonnage is very low de

This 80,000 tons actually included 517 ships. Ship tonnage was very low.

In our NLP applications, we use a level of linguistic structure driven by the argument-taking properties of predicates and composed monotonically up to the sentence level. The resulting Lexical Conceptual Structures (LCS) (Jackendoff, 1983), is a language-neutral representation of the situation (event or state), suitable for use as an interlingua, e.g., for machine translation. The LCS represents predicate argument structure abstracted away from language-specific properties of semantics and syntax. The

* We gratefully acknowledge DOD support for this work through contract MDA904-96-R-0738

primitives of the interlingua provide for monotonic composition that captures both conceptual and syntactic generalities (Dorr et al., 1993) among languages.¹ The strength of the representation derives from the cross-linguistic regularities in the lexical semantics encoded in the LCS. The syntactic hierarchy (subject, object, oblique) is mirrored in the LCS hierarchy: for example THEMES are arguments of the LCS predicate, and AGENTS are arguments of the theme-predicate composition. Syntactic divergences (whether the object precedes or follows the verb, for example) are represented in language specific linearization rules; lexical divergences (whether the location argument is encoded directly in the verb, e.g. the English verb *pocket* or must be saturated by an external argument) are stated in terms of the pieces of LCS structure in the lexicon. Sentential representations derive from saturating the arguments required by the predicates in the sentence.

LCS representations also include temporal information, where available in the source language: recent revisions include, for example (Dorr and Olsen, 1997a) standardizing LCS representations for the aspectual (un)boundedness ((A)TELICITY) of events, either lexically or sententially represented. Although at present the LCS encodes no supra-sentential discourse relations, we show how the lexical aspect information may be used to generate discourse coherence in temporal structure. Relations between clauses as constrained by temporal reference has been examined in an LCS framework by Dorr and Gaasterland (Dorr and Gaasterland, 1995). They explore how temporal connectives are constrained in interpretation, based on the tense of the clauses they connect. While overt temporal connectives are helpful when they appear, our corpus contains many sentences with neither tense markers nor tense connectives. We must therefore look to a new source of information. We rely on the lexical information of the verbs within a sentence to generate both tense and temporal connectives.

Straightforward LCS analysis of many of the multi-clause sentences in our corpus leads to violations of the wellformedness conditions, which prevent structures with events or states directly modifying other events or states. LCS, as previously conceived, prohibits an event or state from standing in a modifier relationship to another event or state, without mediation of a path or position (i.e., as lexically realized by a preposition). This restriction reflects the insight that (at least in English) when events and states modify each other, the modification is either implicit, with the relevant events and states in separate sentences (and hence separate LCSs), as in the

first sentence below, or explicit in a single sentence, as in the second sentence below. Implicit event-state modification (sentence 3) is prohibited.

- Wade bought a car. He needed a way to get to work.
- Wade bought a car because he needed a way to get to work.
- * Wade bought a car he needed a way to get to work.

It is exactly these third type that are permitted in standard Chinese and robustly attested in our data. If the LCS is to be truly an interlingua, we must extend the representation to allow these kinds of sentences to be processed. One possibility is to posit an implicit position connecting the situations described by the multiple clauses. In the source language analysis phase, this would amount to positing a disjunction of all possible position relations implicitly realizable in this language. Another option is to relax the wellformedness constraints to allow an event to directly modify another event. This not only fails to recognize the regularities we see in English (and other language) LCS structures, for Chinese it merely pushes the problem back one step, as the set of implicitly realizable relations may vary from language to language and may result in some ungrammatical or misleading translations. The second option can be augmented, however, by factoring out of the interlingua (and into the generation code) language-specific principles for generating connectives using information in the LCS-structure, proper. For the present, this is the approach we take, using lexical aspectual information, as read from the LCS structure, to generate appropriate temporal relations.

Therefore not only tense, but inter-sentential discourse relations must be considered when generating English from Chinese, even at the sentence level. We report on a project to generate both temporal and discourse relations using the LCS representation. In particular, we focus on the encoding of the lexical aspect feature TELICITY and its complement ATELICITY to generate past and present tense, and corresponding temporal relations for modifying clauses within sentences. While we cannot at present directly capture discourse relations, we can garner aspectual class from LCS verb classification, which in turn can be used to predict the appropriate tense for translations of Chinese verbs into English.

2 Use of Aspect to Provide Temporal Information

We begin with a discussion of aspectual features of sentences, and how this information can be used to provide information about the time of the situations

¹LCS representations in our system have been created for Korean, Spanish and Arabic, as well as for English and Chinese.

presented in a sentence. Such information can be used to help provide clues as to both tense and relationships (and cue words) between connected situations. Aspectual features can be divided into grammatical aspect, which is indicated by lexical or morphological markers in a sentence, and lexical aspect, which is inherent in the meanings of words.

2.1 Grammatical aspect

Grammatical aspect provides a viewpoint on situation (event or state) structure (Smith, 1997). Since imperfective aspect, such as the English PROGRESSIVE construction *be VERB-ing*, views a situation from within, it is often associated with present or contemporaneous time reference. On the other hand, perfective aspect, such as the English *have VERB-ed*, views a situation as a whole; it is therefore often associated with past time reference ((Comrie, 1976; Olsen, 1997; Smith, 1997) cf. (Chu, 1998)). The temporal relations are tendencies, rather than an absolute correlation: although the perfective is found more frequently in past tenses (Comrie, 1976), both imperfective and perfective co-occur in some language with past, present, and future tense.

In some cases, an English verb will specify tense and/or aspect for a complement. For example, *continue* requires either an infinitive (3)a or progressive complement (3)b (and subject drop), while other verbs like *say* do not place such restrictions (3)c,d.

- (3) a. Wolfe continued to publicize the baseless criticism on various occasions
- b. Wolfe continued publicizing the baseless criticism on various occasions
- c. Wolfe continued publicizing the baseless criticism on various occasions
- d. He said the asia-pacific region already became a focal point region
- e. He said the asia-pacific region already is becoming a focal point region

2.2 Lexical aspect

While grammatical aspect and overt temporal cues are clearly helpful in translation, there are many cases in our corpus in which such cues are not present. These are the hard cases, where we must infer tense or grammatical aspectual marking in the target language from a source that looks like it provides no overt cues. We will show however, that Chinese does provide implicit cues through its lexical aspect classes. First, we review what lexical aspect is.

Lexical aspect refers to the type of situation denoted by the verb, alone or combined with other sentential constituents. Verbs are assigned to lexical aspect classes based on their behavior in a variety of

syntactic and semantic frames that focus on three aspectual features: telicity, dynamicity and durativity. We focus on telicity, also known as BOUNDEDNESS. Verbs that are telic have an inherent end: *winning*, for example, ends with the finish line. Verbs that are atelic do not name their end: *running* could end with a distance *run a mile* or an endpoint *run to the store*, for example. Olsen (Olsen, 1997) proposed that aspectual interpretation be derived through monotonic composition of marked privative features [+/ \emptyset dynamic], [+/ \emptyset durative] and [+/ \emptyset telic], as shown in Table 1 (Olsen, 1997, pp. 32–33).

With privative features, other sentential constituents can add to features provided by the verb but not remove them. On this analysis, the [+durative, +dynamic] features of *run* propagate to the sentence level in *run to the store*; the [+telic] feature is added by the NP or PP, yielding an accomplishment interpretation. The feature specification of this compositionally derived accomplishment is therefore identical to that of a sentence containing a telic accomplishment verb, such as *destroy*.

According to many researchers, knowledge of lexical aspect—how verbs denote situations as developing or holding in time—may be used to interpret event sequences in discourse (Dowty, 1986; Moens and Steedman, 1988; Passoneau, 1988). In particular, Dowty suggests that, absent other cues, a telic event is interpreted as completed before the next event or state, as with *ran into the room* in 4a; in contrast, atelic situations, such as *run*, *was hungry* in 4b and 4c, are interpreted as contemporaneous with the following situations: *fell* and *made a pizza*, respectively.

- (4) a. Mary ran into the room. She turned on her walkman.
- b. Mary ran. She turned on her walkman.
- c. Mary was hungry. She made a pizza.

Smith similarly suggests that in English all past events are interpreted as telic (Smith, 1997) (but cf. (Olsen, 1997)).

Also, these tendencies are heuristic, and not absolute, as shown by the examples in (5). While we get the expected prediction that the jumping occurs after the explosion in (5)(a), we get the reverse prediction in (5)(b). Other factors such as consequences of described situations, discourse context, and stereotypical causal relationships also play a role.

- (5) a. The building exploded. Mary jumped.
- b. The building exploded. Chunks of concrete flew everywhere.

| Aspectual Class | Telic | Dynamic | Durative | Examples |
|-----------------|-------|---------|----------|--------------------|
| State | | | + | <i>know, have</i> |
| Activity | | + | + | <i>run, pain</i> |
| Accomplishment | + | + | + | <i>destroy</i> |
| Achievement | + | + | | <i>notice, win</i> |

Table 1: Lexical Aspect Features

3 Aspect in Lexical Conceptual Structure

Our implementation of Lexical Conceptual Structure (Dowty, 1979; Guerssel et al., 1985)—an augmented form of (Jackendoff, 1983; Jackendoff, 1990)—permits lexical aspect information to be read directly off the lexical entries for individual verbs, as well as composed representations for sentences, using uniform processes and representations. The LCS framework consists of primitives (GO, BE, STAY, etc.), types (Event, State, Path, etc.) and fields (Loc(ational), Temp(oral), Poss(essional), Ident(ificational), Perc(eptual), etc.).

We adopt a refinement of the LCS representation, incorporating meaning components from the linguistically motivated notion of *lexical semantic template* (LST), based on lexical aspect classes, as defined in the work of Levin and Rappaport Hovav (Levin and Rappaport Hovav, 1995; Rappaport Hovav and Levin, 1995). Verbs that appear in multiple aspectual frames appear in multiple pairings between constants (representing the idiosyncratic meaning of the verb) and structures (the aspectual class). Since the aspectual templates may be realized in a variety of ways, other aspects of the structural meaning contribute to differentiating the verbs from each other. Our current database contains some 400 classes, based on an initial representation of the 213 classes in (Levin, 1993). Our current working lexicon includes about 10,000 English verbs and 18,000 Chinese verbs spread out into these classes.

Telic verbs (and sentences) contain certain types of Paths, or a constant, represented by `!!`, filled by the verb constant, in the right most leaf-node argument. Some examples are shown below:

```
depart (go loc (* thing 2)
        (away_from loc (thing 2)
          (at loc (thing 2)
            (* thing 4)))
        (!!+ingly 26))

insert (cause (* thing 1)
        (go loc (* thing 2)
          ((* toward 5) loc (thing 2)
```

```
([at] loc (thing 2)
          (thing 6)))
 (!!+ingly 26))
```

Each of these telic verbs has a potential counterpart with an atelic verb plus the requisite path. *Depart*, for example, corresponds to *move away*, or something similar in another language.

We therefore identify telic sentences by the algorithm, formally specified in in Figure 1 (cf. (Dorr and Olsen, 1997b) [156]).

Given an LCS representation L:

1. Initialize: T(L):={∅T}, D(L):={∅R}, R(L):={∅D}
2. If Top node of L ∈ {CAUSE, LET, GO}
 - Then T(L):={+T}
 - If Top node of L ∈ {CAUSE, LET}
 - Then D(L):={+D}, R(L):={+R}
 - If Top node of L ∈ {GO}
 - Then D(L):={+D}
3. If Top node of L ∈ {ACT, BE, STAY}
 - Then If Internal node of L ∈ {TO, TOWARD, FOR_{Temp}}
 - Then T(L):={+T}
 - If Top node of L ∈ {BE, STAY}
 - Then R(L):={+R}
 - If Top node of L ∈ {ACT}
 - Then set D(L):={+D}, R(L):={+R}
4. Return T(L), D(L), R(L).

Figure 1: Algorithm for Aspectual Feature Determination

This algorithm applies to the structural primitives of the interlingua structure rather than actual verbs in source or target language. The first step initialized the aspectual values as unspecified: atelic [-T], stative (not event: [-D]), and adurative [-R]. First the top node is examined for primitives that indicate telicity: if the top node is CAUSE, LET, GO, telicity is set to [+T], as with the verbs *break*, *destroy*, for example. (The node is further checked for dynamicity [+D] and durativity [+R] indicators, not in focus in this paper.) If the top node is not a telic indicator (i.e., the verb is a basically atelic predicate such as *love* or *run*, telicity may still be indicated by the presence of complement nodes of particular types: e.g. a goal phrase (to primitive) in the case of *run*. The same algorithm may be used to determine telicity in either individual verbal entries (*break* but

not *run*) or composed sentences (*John ran to the store but not John ran*).

Similar mismatches of telicity between representations of particular predicates can occur between languages, although there is remarkable agreement as to the set of templates that verbs with related meanings will fit into (Olsen et al., 1998). In the Chinese-English interlingual system we describe, the Chinese is first mapped into the LCS, a language-independent representation, from which the target-language sentence is generated. Since telicity (and other aspects of event structure) are uniformly represented at the lexical and the sentential level, telicity mismatches between verbs of different languages may then be compensated for by combining verbs with other components.

4 Predictions

Based on (Dowty, 1986) and others, as discussed above, we predict that sentences that have a telic LCS will better translate into English as the past tense, and those that lack telic identifiers will translate as present tense. Moreover, we predict that verbs in the main clause that are telic, will be past with respect to their subordinates (X then Y). Verbs in the main clause that are atelic we predict will temporally overlap (X while Y).

5 Implementation

LCSes are used as the interlingua for our machine translation efforts. Following the principles in (Dorr, 1993), lexical information and constraints on well-formed LCSes are used to *compose* an LCS for a complete sentence from a sentence parse in a source language. This composed LCS (CLCS) is then used as the starting points for generation into the target language, using lexical information and constraints for the target language.

The generation component consists of the following subcomponents:

Decomposition and lexical selection First, primitive LCSes for words in the target language are matched against CLCSes, and tree structures of covering words are selected. Ambiguity in the input and analysis represented in the CLCS is maintained (insofar as it is possible to realize particular readings using the target language lexicon), and new ambiguities are introduced when there are different ways of realizing a CLCS in the target language.

AMR Construction This tree structure is then translated into a representation using the Augmented Meaning Representation (AMR) syntax of instances and hierarchical relations (Langkilde and Knight, 1998a); however the relations include information present in the CLCS

and LCSes for target language words, including theta roles, LCS type, and associated features.

Realization The AMR structure is then linearized, as described in (Dorr et al., 1998), and morphological realization is performed. The result is a lattice of possible realizations, representing both the preserved ambiguity from previous processing phases and multiple ways of linearizing the sentence.

Extraction The final stage uses a statistical bigram extractor to pick an approximation of the most fluent realization (Langkilde and Knight, 1998b).

While there are several possible ways to address the tense and discourse connective issues mentioned above, such as modifying the LCS primitive elements and/or the composition of the LCS from the source language, we instead have been experimenting for the moment with solutions implemented within the generation component. The only extensions to the LCS language have been loosening of the constraint against direct modification of states and events by other states and events (thus allowing composed LCSes to be formed from Chinese with these structures, but creating a challenge for fluent generation into English), and a few added features to cover some of the discourse markers that are present. We are able to calculate telicity of a CLCS, using the algorithm in Figure 1 and encode this information as a binary **telic** feature in the Augmented Meaning Representation (AMR).

The realization algorithm has been augmented with the rules in (6)

(6) a. If there is no tense feature, use telicity to determine the tense:

:telic + → :tense past
:telic - → :tense present

b. In an event or state directly modifying another event or state, if there is no other clausal connective (coming from a subordinating conjunction or post-position in the original), then use telicity to pick a connective expressing assumed temporal relation:

:telic + → :sconj then
:telic - → :sconj while

6 The Corpus

We have applied this machine translation system to a corpus of Chinese newspaper text from *Xinhua* and other sources, primarily in the economics domain. The genre is roughly comparable to the American

Wall Street Journal. Chinese newspaper genre differs from other Chinese textual sources, in a number of ways, including:

- more complex sentence structure
- more extensive use of acronyms
- less use of Classical Chinese
- more representative grammar
- more constrained vocabulary (limited lexicon)
- abbreviations are used extensively in Chinese newspaper headlines

However, the presence of multiple events and states in a single sentence, without explicit modification is characteristic of written Chinese in general. In the 80-sentence corpus under consideration, the sentence structure is complex and stylized; with an average of 20 words per sentence. Many sentences, such as (1) and (2), have multiple clauses that are not in a direct complement relationship or indicated with explicit connective words.

7 Ground Truth

To evaluate the extent to which our predictions result in an improvement in translation, we have used a database of human translations of the sentences in our corpus as the ground truth, or gold standard. One of the translators is included among our authors.

The ground truth data was created to provide a fluid human translation of the text early in our system development. It therefore includes many complex tenses and multiple sentences combined, both currently beyond the state of our system. Thus, two of the authors and an additional researcher also created a database of temporal relations among the clauses in the sentences that produced illegal event/state modifications. This was used to test predictions of temporal relationships indicated by telicity. In evaluating our results, we concentrate on how well the system did at matching past and present, and on the appropriateness of temporal connectives generated.

8 Results

We have applied the rules in (6) in generating 80 sentences in the corpus (starting from often ambiguous CLCS analyses). Evaluation is still tricky, since, in many cases, the interlingua analysis is incorrect or ambiguous in ways that affect the appropriateness of the generated translation.

8.1 Tense

As mentioned above, evaluation can be very difficult in a number of cases. Concerning tense, our "gold standard" is the set of human translations,

| | | generated tense | |
|-------------------|---------|-----------------|---------|
| | | past | present |
| human translation | past | 34 | 17 |
| | present | 17 | 27 |

Table 2: Preliminary Tense Results

previously constructed for these sentences. In many cases, there is nothing overt in the sentence which would specify tense, so a mismatch might not actually be "wrong". Also, there are a number of sentences which were not directly applicable for comparison, such as when the human translator chose a different syntactic structure or a complex tense. The newspaper articles were divided into 80 sentences. Since some of these sentences were conjunctions, this yielded 99 tensed main verbs. These verbs either appeared in simple present, past, present or past perfect ('has or had verb+ed), present or past imperfective (is verb+ing, was verb+ing) and their corresponding passive (is being kicked, was being kicked, have been kicked) forms. For cases like the present perfect ('has kicked), we noted the intended meaning (e.g past activity) expressed by the verb as well as the verb's actual present perfective form. We scored the form as correct if the system translated a present perfective with past tense meaning as a simple past or present perfective. There were 10 instances where a verb in the human translation had no corresponding verb in the machine translation, either due to incorrect omission or correct substitution of the corresponding nominalization. We excluded these forms from consideration. If the system fails to supply a verb for independent reasons, our system clearly can't mark it with tense. The results of our evaluation are summarized in Table 2.

These results definitely improve over our previous heuristic, which was to always use past tense (assuming this to be the default mode for newspaper article reporting). Results are also better than always picking present tense. These results seem to indicate that atelicity is a fairly good cue for present tense. We also note that 8 out of the 14 cases where the human translation used the present tense while the system used past tense are headlines. Headlines are written using the historical present in English ("Man bites Dog"). These sentences would not be incorrectly translated in the past ("The Man Bit the Dog") Therefore, a fairer judgement might leave only remaining 6 incorrect cases in this cell. Using atelicity as a cue for the present yields correct results approximately 65 incorrect results 35 worst case results because they do not take into account presence or absence of the grammatical perfective and progressive markers referred to in the introduction.

8.2 Relationship between clauses

Results are more preliminary for the clausal connectives. Of the 80 sentences, 35 of them are flagged as (possibly) containing events or states directly modifying other events or states. However, of this number, some actually do have lexical connectives represented as featural rather than structural elements in the LCS, and can be straightforwardly realized using translated English connectives such as *since*, *after*, and *if-then*. Other apparently "modifying" events or states should be treated as a complement relationship (at least according to the preferred reading in ambiguous cases), but are incorrectly analyzed as being in a non-complement relationship, or have other structural problems rendering the interlingua representation and English output not directly related to the original clause structure.

Of the remaining clear cases, six *while* relationships were generated according to our heuristics, indicating cotemporaneity of main and modifying situation, e.g. (7)a,b, in the automated translations of (1) and (2), respectively. None were inappropriate. Of the cases where *then* was generated, indicating sequential events, there were four cases in which this was appropriate, and three cases in which the situations really should have been cotemporaneous. While these numbers are small, this preliminary data seems to suggest again that atelicity is a good cue for cotemporality, while telicity is not a sufficient cue.

- (7) a. Before 1965, China altogether only have the ability shipbuilding about 300 thousand tons , while the annual output is 80 thousand tons .
- b. this 80 thousand tons actually includes 517 ships, while the ship tonnage is very low .

9 Conclusions

We therefore conclude that lexical aspect can serve as a valuable heuristic for suggesting tense, in the absence of tense and other temporal markers. We anticipate incorporation of grammatical aspect information to improve our temporal representation further. In addition, lexical aspect, as represented by the interlingual LCS structure, can serve as the foundation for language specific heuristics. Furthermore, the lexical aspect represented in the LCS can help to provide the beginnings of cross-sentential discourse information. We have suggested applications in the temporal domain *while*, *then*. Causality is another possible domain in which relevant pieces encoded in sentence-level LCS structures could be used to provide links between LCSes/sentences. Thus, the interlingual representation may be used to provide not only shared semantic and syntactic structure, but also the building blocks for language-specific heuristics for mismatches between languages.

10 Future Research

There are a number of other directions we intend to pursue in extending this work. First, we will evaluate the role of the grammatical aspect markers mentioned above, in combination with the telicity features. Second, we will also examine the role of the nature of the modifying situation. Third, we will incorporate other lexical information present in the sentence, including adverbial cue words (e.g. *now*, *already* and specific dates that have time-related information, and distinguishing reported speech from other sentences. Finally, as mentioned, these results do not take embedded verbs or verbs in adjunct clauses into account. Many adjunct and embedded clauses are tenseless, making evaluation more difficult. For example, is *The President believed China to be a threat* equivalent to *The president believed China is a threat*).

References

- Chauncey C. Chu. 1998. *A Discourse Grammar of Mandarin Chinese*. Peter Lang Publishing, Inc., New York, NY.
- Bernard Comrie. 1976. *Aspect*. Cambridge University Press, Cambridge, MA.
- Bonnie J. Dorr and Terry Gaasterland. 1995. Selecting Tense, Aspect, and Connecting Words in Language Generation. In *Proceedings of IJCAI-95*, Montreal, Canada.
- Bonnie J. Dorr and Mari Broman Olsen. 1997a. Aspectual Modifications to a LCS Database for NLP Applications. Technical Report LAMP TR 007, UMIACS TR 97-23, CS TR 3763, University of Maryland, College Park, MD.
- Bonnie J. Dorr and Mari Broman Olsen. 1997b. Deriving Verbal and Compositional Lexical Aspect for NLP Applications. In *Proceedings of the 35th Annual Meeting of the Association for Computational Linguistics (ACL-97)*, pages 151-158, Madrid, Spain, July 7-12.
- Bonnie J. Dorr, James Hendler, Scott Blanksteen, and Barrie Migdaloff. 1993. Use of Lexical Conceptual Structure for Intelligent Tutoring. Technical Report UMIACS TR 93-108, CS TR 3161, University of Maryland.
- Bonnie J. Dorr, Nizar Habash, and David Traum. 1998. A Thematic Hierarchy for Efficient Generation from Lexical-Conceptual Structure. In *Proceedings of the Third Conference of the Association for Machine Translation in the Americas, AMTA-98, in Lecture Notes in Artificial Intelligence, 1529*, pages 333-343, Langhorne, PA, October 28-31.
- Bonnie J. Dorr. 1993. *Machine Translation: A View from the Lexicon*. The MIT Press, Cambridge, MA.

- David Dowty. 1979. *Word Meaning in Montague Grammar*. Reidel, Dordrecht.
- David Dowty. 1986. The Effects of Aspectual Class on the Temporal Structure of Discourse: Semantics or Pragmatics? *Linguistics and Philosophy*, 9:37-61.
- Mohamed Guerssel, Kenneth Hale, Mary Laughren, Beth Levin, and Josie White Eagle. 1985. A Cross-linguistic Study of Transitivity Alternations. In W. H. Eilfort, P. D. Kroeber, and K. L. Peterson, editors, *Papers from the Parasession on Causatives and Agentivity at the Twenty-First Regional Meeting, CLS 21, Part 2*, pages 48-63. The Chicago Linguistic Society, Chicago, IL, April.
- Ray Jackendoff. 1983. *Semantics and Cognition*. The MIT Press, Cambridge, MA.
- Ray Jackendoff. 1990. *Semantic Structures*. The MIT Press, Cambridge, MA.
- Irene Langkilde and Kevin Knight. 1998a. Generation that Exploits Corpus-Based Statistical Knowledge. In *Proceedings of COLING-ACL '98*, pages 704-710.
- Irene Langkilde and Kevin Knight. 1998b. The Practical Value of N-Grams in Generation. In *International Natural Language Generation Workshop*.
- Beth Levin and Malka Rappaport Hovav. 1995. *Unaccusativity: At the Syntax-Lexical Semantics Interface*. The MIT Press, Cambridge, MA. LI Monograph 26.
- Beth Levin. 1993. *English Verb Classes and Alternations: A Preliminary Investigation*. University of Chicago Press, Chicago, IL.
- Charles Li and Sandra Thompson. 1981. *Mandarin Chinese: A functional reference grammar*. University of California Press, Berkeley, CA.
- Marc Moens and Mark Steedman. 1988. Temporal Ontology and Temporal Reference. *Computational Linguistics: Special Issue on Tense and Aspect*, 14(2):15-28.
- Mari Broman Olsen, Bonnie J. Dorr, and Scott C. Thomas. 1998. Enhancing Automatic Acquisition of Thematic Structure in a Large-Scale Lexicon for Mandarin Chinese. In *Proceedings of the Third Conference of the Association for Machine Translation in the Americas, AMTA-98, in Lecture Notes in Artificial Intelligence, 1529*, pages 41-50, Langhorne, PA, October 28-31.
- Mari Broman Olsen. 1997. *A Semantic and Pragmatic Model of Lexical and Grammatical Aspect*. Garland, New York.
- Rebecca Passoneau. 1988. A Computational Model of the Semantics of Tense and Aspect. *Computational Linguistics: Special Issue on Tense and Aspect*, 14(2):44-60.
- Malka Rappaport Hovav and Beth Levin. 1995. The Elasticity of Verb Meaning. In *Processes in Argument Structure*, pages 1-13, Germany. Sfs-Report-06-95, Seminar für Sprachwissenschaft, Eberhard-Karls-Universität Tübingen, Tübingen.
- Carlota Smith. 1997. *The parameter of aspect*. Kluwer, Dordrecht, 2nd edition.