Automatic Semantic Annotation of Texts

Fernando Gomez
School of EECS
University of Central Florida, Orlando, FL 32816
gomez@cs.ucf.edu

Abstract

This paper explains the automatic semantic interpretation of sentences. These sentences were selected for testing the verb predicates that we have been defining for WordNet 1.6 verb classes, and the algorithm that uses them. Semantic interpretation is performed for determining the meaning of the verb, or, verb predicate, its semantic roles, adjuncts, the attachment of postverbal prepositional phrases, and also for a limited number of nominalizations. A comparison to other approaches to semantic annotation is provided.

1 Introduction

We explain the automatic semantic annotation of sentences using our semantic interpreter (Gomez, 2001; Gomez, 2004a). We have produced a corpus of 2000 sentences, taken from the *The World Book Encyclopedia*, World Book, Inc. (Chicago). One thousand of those sentences can be downloaded from our homepage (www.cs.ucf.edu/ğomez). The remaining one thousand sentences can be obtained upon requesting them from the author of this paper. Some of the sentences have undergone editing so that they could be processed by our system. These sentences were generated for validating some of the verb predicates (Gomez, 1998) that we have been defining for WordNet 1.6 (henceforth, WN) verb classes (Fellbaum, 1998). We think that these sentences provide a panorama of the predicates and its semantic roles for a relatively large class of verbs, and also, that these interpreted sentences could be of value for machine-learning researchers that need interpreted corpora, and for the NLP community at large. By semantic interpretation, we mean the determination of the meaning of the verb, or, predicate, its semantic roles, adjuncts, and the attachment of prepositional phrases. The senses of nouns are also resolved (see below for details). But, compound nouns are not resolved, unless the pair of nouns exist in the WN lexicon. We have also provided semantic interpretation for some of the deverbal nominalizations in the texts if the nominalization is in our database of nominalizations, which is incomplete. No resolution of anaphora is done, which is beyond the scope of this work.

The parser that builds the input for the semantic interpreter does not resolve structural ambiguity, but does identify verb complements including clausal complements. Thus, the attachment of PPs and any postverbal modifiers are the tasks of the semantic interpreter. In order to test the predicates, it is not necessary to parse complex sentences containing relativization, coordination and/or subordination, free relatives, etc. The only thing needed is to pass to the semantic interpreter all the complements (including clausal complements) and postverbal modifiers of the verb in the clause. However in our data set, we have processed sentences containing relativization, coordination, subordination and adverbial clauses in order to provide a rich set of syntactic and semantic structures.

2 Output of the Semantic Interpreter

The output of the semantic interpreter is very easy to read. The sentence is listed first followed by the semantic interpretation for that sentence, which, in turn, is followed by the hierarchy of all super-predicates for the verb predicates in the sentence. The grammatical relation is listed first followed by the semantic role for that grammatical relation in the form $< (semantic - role) >$. Hence, the pair <> identifies semantic roles, which appear always

```
(CLAUSE CL1 (P "Shotgun shells are filled with small pellets")
 (SUBJ ((NOUN SHOTGUN) (NOUN SHELLS)) ((SHELL1 SHOTGUN_SHELL1)) <GOAL>)
 (VERB FILLED ((AUX (ARE)) (MAIN-VERB FILL FILLED) (VOICE PASSIVE))
       FILL-SOMETHING (FILL1 FILL2) SUPPORTED BY 2 SRS)
 (PREP WITH  (PREP-NP ((ADJ SMALL) (NOUN PELLETS))
              ((SHOT6 PELLET2)) <THEME>)
               (ATTACH VERB CONFIDENCE STRONG))
)
FILL-SOMETHING => FILL-OR-LOAD => PUT => CAUSE-TO-CHANGE-LOCATION => ACTION
```

Figure 1: Output for the sentence "Shotgun shells are filled with small pellets."

to the right of the grammatical relation. The algorithm selects the predicate that maximizes semantic roles, or explains the most grammatical relations (see (Gomez, 2001) for details). The program prints the predicate followed by the WN verbs senses and the phrase "supported by 'a-number' SRs," (where 'a-number' stands for a cardinal number), which indicates the number of grammatical relations (SRs) that have realized semantic roles. If the output for that sentence has more semantic roles than the number listed in the phrase "supported by," those semantic roles are adjuncts (Gomez, 2001).

Hierarchy of Predicates

After printing the semantic interpretation for the sentence, the hierarchy for the predicates in the sentence is printed. This hierarchy and the semantic roles for the predicate determine the meaning of the predicate. In the first file in our corpus (file f50-1.l), if a hierarchy has been already given for a predicate that hierarchy is not listed again. However in the remainder of the files, the hierarchy for a predicate is listed every time in which the predicate is identified in a sentence.

For some predicates, the algorithm prints the predicate concatenated with "-I" or "-II" Thus, we may have the hierarchy

```
LEAVE-A-PLACE-I
  LEAVE-A-PLACE
    CHANGE-LOCATION
      ACTION
       ROOT
```

The predicate LEAVE-A-PLACE-I is identical to LEAVE-A-PLACE, which is listed as its immediate super-predicate. Thus, the predicates *predicate-name-I* and *predicate-name-II* are identical to *predicate-name*. The use of "-I" and "-II" became necessary because of the semi-automatic definition of some predicates. In our current implementation, we have eliminated the use of the suffixes [-I] and [-II], and predicates are listed without any of these suffixes. Hence, the hierarchy above will begin with *leave-a-place*.

3 WordNet Verb Senses and Noun Senses

Sometimes more than one WN verb sense is given for a predicate. For instance, for the predicate *travel-change-location* given for the verb "travel" in the sentence "She traveled to Spain," the following WN senses of "travel" are listed (travel1 travel4 travel5). This means that all WN senses of "travel" in that list are coalesced into the predicate *travel-change-location* (Gomez, 2001). Hence, all these WN senses of the verb have the same meaning.

In some cases, the algorithm provides the sense of the first verb listed in the WN synset list. For instance, in the sentence "She carried the books to the store," the algorithm provides the sense transport2 for the predicate *carry-something*. Transport2 is *carry1* in WN. The sense carry1 can be easily obtained from transport2, something we have not done because of the many things that needed implementation. Likewise, for "pass a law, bill, etc." the algorithm may give you "legislate1," which is pass3. If the verb is immediately followed by a question mark, the sense of the verb does not exist in WN, or it is unclear which WN verb sense the verb in the sentence may stand for.

Noun Senses

In some instances, the algorithm may provide more than one sense for a noun. To the contrary

to verb senses, this list is a ranked list in which the first noun sense listed by the algorithm is the preferred one. For instance, in the sentence "The star remains stable," two senses are provided for "star" as follows: (CELESTIAL-BODY STAR1 STAR3). Both senses of "star" have *celestial body* as their hypernym. The difference between "star1" and "star3" is that "start3" is visible to the naked eye. The algorithm is unable to distinguish between these two senses of "star" and lists both senses. In the sentence "Cats climb trees," the algorithm is unable to determine if "cat" refers to a domestic cat (cat1), or to a big cat (cat6). Thus, the senses for "cat" are listed as: (ANIMAL CAT1 CAT6).

The ontological categories without numbers correspond to our ontological categories, which are the result of some changes and additions to WN 1.6 upper-level ontology (Miller, 1998) (see (Gomez, 2004b) for details).

4 Semantic Roles and Predicate Classes

Our approach to semantic roles is better described in the discussion in (Gomez, 1998). Briefly, it is as follows. Predicates are organized in an inheritance hierarchy in a similar way as ontological categories for nouns. There is one *generic* predicate, or super-predicate, for each class of predicates. Semantic roles depend on this generic predicate which determines their number and their meaning. Thus, the generic predicate **communicate** that contains *transfer-of-information* as its major subpredicate, has the following semantic roles: an *agent* (the person, or animal that communicates), a *theme*, the thing communicated, a *recipient*, the person or animal that receives the thing being communicated, the *form-or-medium-of-expression* in which the communication takes place, e.g. "book" in "This book explains algebra to children," and the *topic*. Some of the subpredicates of *communicate* may have some semantic roles that are specific to these predicates. For instance, *thank-somebody* has one semantic role realized by the PP [for + NP] which is specific to this predicate. We have defined 243 subpredicates of *communicate*.

The predicate class **change-location** stands for a person or animal changing location from a place to another. In this predicate, the *agent* and *theme* are the same. Other roles are: *to-loc*, the destination of the agent, *from-loc*, the place from where the agent changed location, the *instrument*, the means used to change location, the *distance*, which is the distance traveled, and *at-speed* the speed at which the change of location takes place.

The generic predicate **cause-to-change-location** expresses a cause of change of location of something other than the agent; although the agent may have also changed location. In "The squirrel carried the pecans to the nest" the agent has also been moved, but the primary event is the fact that the squirrel caused a change of location of the pecans. In a sentence such as "The moon circles the earth," "the moon" is the *theme* and the *inanimate-cause* is unknown. The semantic roles for this generic predicate and its subpredicates are: *agent*, *theme*, *goal*, *source*, *distance*, *inanimate-cause* and *instrument*. Upon completion of the event, the entity expressed by the *theme* is at the *goal*, while the entity expressed by the *theme* was at the *source*. Thus, in "Kelly filled the tank with water from the river," "the tank" is the *goal*, "the water" is the *theme*" and "the river" is the source. The *inanimate-cause* is any entity other than a human-agent or an animal used in a causal way, e.g. "The wind pulled the trees from the ground," "The war caused much suffering," "His policies caused many revolts" etc.

Stative Predicates form several classes. Many of them belong to the classes of verb senses listed under be1, be2 and be3 in WN. We have used the term *description* for referring to the parent predicate of all these predicates. Some examples of generic predicates within this class are *be-at-a-place*, *include*, etc. We have used the term *thing-described* for referring to the entity of which something is being predicated. Thus, in the sentence "The book is on the table," the "book" is the *thing-described* and "on the table" is the *at-loc*. In "The table is red" "the table" is the *thing-described* and "red" is the *attribute-described* of the table. We have also used *attribute-described* to refer to the semantic roles realized by the direct objects of stative verbs. For instance, in "This book includes many references," "many references" is

the *attribute-described*.

The roles for the generic predicate **cause-change-of-state** are: *agent*, *theme*, *beginning-state*, *ending-state*, *inanimate-cause* and *instrument*. In these predicates, the *theme* has undergone a change of state, not just a change of location, or a transfer. There is a big difference between the *ending-state* of this predicate and the *goal* of the predicate *cause-to-change-location*. In "He broke the cup into small pieces," the cup has undergone a change of state, and does not exist any longer in the *ending-state*. In "She shaped the metal into a beautiful statue," the metal has been transformed into a statue in the *ending-state*. There are major subpredicates of the predicate *cause-to-change-state* for which we have defined over 600 subpredicates. Some of the major sub-predicates of the predicate *cause-to-change-state* have little relation between each other. They range from *cause-change-of-integrity* to *cause-a-change-of-human-belief*, e.g., "He converted the Irish to Catholicism" in which the *ending-state* expresses a change of belief by the *theme*. These major sub-predicates are well contained within themselves, and they inherit little from their super-predicate *cause-to-change-state*. We could have defined semantic roles that are specific to them, and we have done so in some cases. These semantic roles are more informative than the semantic roles of their generic predicates. For instance, for the sub-predicate *separate-from-somebody*, separate8 in WN ("discontinue an association or relation; go different ways"), the *beginning-state* is expressed as *from-person*, meaning the person or social group one separates from.

Hierarchy of Semantic Roles

This results into a hierarchy of semantic roles. Thus, one has *from-person* => *beginning-state*. (We have used => to indicate the super-role) Examples of other roles and their super-roles found in the corpus are:

in-court-tribunal => at-loc.

at-inquiry-investigation => *at-activity*.

for-performance => *purpose*.

because-of-recipient-acts => *because-of*.

cure-agent => *inanimate-cause*.

work-place => *at-loc*

from-organization-or-activity => *beginning-state*

One may decide to base the inferences on the super-roles, or on the the more specific roles. Moreover, roles may be defined in terms of predicates and the noun ontological categories, facilitating the inferences associated with them. For instance, a role that is used in the predicate *defend-something* and its sub-predicates (*protect*, *save-something* ...) is *from-attacker*, which is the thing that attacks the entities referred by the *theme* of these predicates. Thus, we have sentences such as "She protected/saved the city from the invaders/the plague." The meaning of *from-attacker* can be easily captured in an axiom expressing the fact that the entities denoted by this role attack the entities denoted by the *theme* of the predicate, resulting into a general inference associated with this role.

5 An Example of Sub-Hierarchy

Sub-hierarchies within a major generic predicate are self-contained, and provide the best way to store the inferences associated with the predicates. In fact, anchoring the inferences into these sub-generic predicates will suffice for most applications. For instance, the predicate *withdraw-from-an-activity-or-organization* => *social-change* => *cause-change-of-state* forms the following sub-hierarchy:

```
WITHDRAW-FROM-AN-ACTIVITY-OR-ORGANIZATION
 GO-ON-RETIREMENT (RETIRE1)
 QUIT-ORGANIZATION-OR-JOB-ACTIVITY
                     (LEAVE8 QUIT2)
  RETIRE-FROM-AN-ACTIVITY (RETIRE2)
  DROP-OUT (DROP_OUT1 DROP_OUT2)
  RESIGN (VACATE1)
    STEP-DOWN-RESIGN (STEP_DOWN1)
  ABDICATE-ORGANIZATION-RESPONSIBILITIES
                       (ABDICATE1)
  GIVE-UP-A-JOB (GIVE_UP7)
  QUIT-ORGANIZATION-OR-JOB(leave_office1)
  DEFECT (DEFECT1)
    DESERT-DEFECT (DEFECT2)
  RENOUNCE-AN-OCCUPATION
            (RENOUNCE1 RENOUNCE2)
   WITHDRAW-A-NAME
```

The predicate refers to a human-agent withdrawing himself/herself/somebody from an activity, or a social group. In a sentence such as "The teacher gave up her job next summer," the *agent* and the *theme* are the same. The meaning

```
(CLAUSE CL21 (P "the teacher gave up her job next summer")
 (SUBJ ((DFART THE) (NOUN TEACHER)) ((EDUCATOR1 TEACHER1)) <AGENT>)
 (VERB GAVE_UP ((MAIN-VERB GIVE_UP GAVE_UP)) GIVE-UP-A-JOB (GIVE_UP7)
       SUPPORTED BY 2 SRS)
 (OBJ ((ADJ HER) (NOUN JOB)) ((OCCUPATION JOB1))
      <FROM-ORGANIZATION-OR-ACTIVITY>)
 (TIME-NP ((ADJ NEXT) (NOUN SUMMER)) ((SEASON2 SUMMER1)) <AT-TIME>))

GIVE-UP-A-JOB => QUIT-ORGANIZATION-OR-JOB => WITHDRAW-FROM-AN-ACTIVITY-OR-ORGANIZATION
   => SOCIAL-CHANGE => CAUSE-CHANGE-OF-STATE => ACTION
```

Figure 2: Output for "The teacher gave up her job next summer."

is that she withdrew herself from her occupation (see figure 2). The activity, or organization are expressed by the role *from-organization-or-activity*. One could split this role into two roles. The reason why they are merged into one is due to simplifying our work in building the entries for the predicate, and also to the semi-automatic construction of the predicates.

6 Noun Senses and Selectional Restrictions

Our corpus of 2000 sentences is the result of validating some of the verb predicates we have constructed. The building of the corpus is not an end in itself but rather a positive side effect from testing the predicate definitions. If the definitions of the predicates are adequate, there is no need to build a corpus because the semantic interpreter must be able to determine the meaning of the verb and its semantic roles in any corpus. Most of the corrections to the output of the semantic interpreter were needed because the algorithm failed to determine the meaning of noun senses. This should not come as a surprise because verb selectional restrictions cannot constrain the sense of the noun in many situations. But, in defining the selectional restrictions for semantic roles we have tried to define them in such way that they narrow down the noun senses to a small number. There are those instances in which the selectional restrictions do narrow the sense of the noun to just one. For instance, if you enter "She filled the mug with beer" the object of the sentence is interpreted as

(OBJ ((DFART THE) (NOUN MUG)) ((DRINKING-VESSEL1 MUG4)) (GOAL))

The system has selected the sense mug4 for "mug" because the *goal* for the verb predicate *fill-something* is:

```
(GOAL (XOR CONTAINER1 PHYSICAL-THING THING)
```

The entry XOR means that only one of the ontological categories listed applies. The system selects the first one that applies. This allows the system to select the sense of mug4 for the word "mug." However, in most situations the system cannot determine the sense of the noun based on the selectional restrictions, but it may prefer some senses over others. For instance the entry for the *theme* of the predicate *put* which has many subpredicates is:

(THEME (AND PHYSICAL-OBJECT PHYSICAL-THING) (OBJ))

The entry AND means that both categories *physical-object* and *physical-thing* apply, but it prefers the first entry listed. Thus, if you enter the sentence "She put the plants in the box," the system outputs for the *theme* of that sentence is:

(OBJ ((DFART THE) (NOUN PLANTS)) ((PLANT PLANT2) (PLANT-1 PLANT1)) (THEME))

The system prefers the sense of plant2 (a life form) over plant1 (a factory) because living forms have *physical-object* as a hypernym in the ontology (see (Gomez, 2004b) for discussion). The selectional restrictions of the *theme* is reordering the WN noun senses and preferring the second sense in WN over the first one. In other cases, the system may exclude some noun senses narrowing the preferred senses to a few ones. For instance, the output for the theme of "She dropped the stalks on the ground" is:

(OBJ ((DFART THE) (NOUN STALKS))

((CHAFF1 STALK1) (PLANT-ORGAN1 STALK2)) (THEME))

Three other senses of "stalk" all of which have *action* as their hypernym have been ruled out by the selectional restrictions of the *theme* of the predicate *drop-something-physical*.

In other cases, the system returns more than one verb predicate for a verb. For the sentence "The plant is widely cultivated for its pungent root," the system returns *cultivate(cultivate1)* and *cultivate-knowledge-wisdom(cultivate3)* in order of preference. The second predicate is also possible because one of the senses of "plant," plant3, has *cognition, knowledge* as its hypernym in WN.

7 Related Work

This work bears relation to Framenet (Fillmore et al., 2003) and Propbank (Kingsbury et al., 2002). But, the major relation is to Framenet. However, there are important differences between our work and Framenet (Fillmore et al., 2003). An obvious difference between Framenet and our approach is that the definitions of our predicates are immediately followed by their validation which consists of running sentences randomly selected from a corpus through our semantic interpreter. A major difference is that the definition of our predicates are linked to selectional restrictions and the grammatical relations that realize them. The selectional restrictions of the predicates are grounded on the WordNet ontology for nouns (Miller, 1998). The definition of the selectional restrictions is what makes it possible for our system to determine the meaning of the verb, or verb predicate, and its semantic roles. To our knowledge, Framenet does not deal with the problem of verb polysemy, which is at the core of the main issue we try to solve. If you consider the sentences "The engine drives the vehicle" vs "The teacher drives the vehicle," or "The nurse drew blood from the doctor" vs "The nurse drew criticism from the doctor" grammatical relations alone cannot help to determine the senses of "drive" and "draw," because both sentences have the same grammatical relations. The only way one can determine the senses of "drive" and "draw" in those examples is by the semantic categories of the entities expressed by the grammatical relations of the sentences; or, in other words, by the selectional restrictions of the verb predicates those verb senses stand for.

A major effort in our work has gone into defining the selectional restrictions of the semantic roles in the verb predicates in such way that those definitions will make it possible the determination of verb meaning, or verb predicate, and its semantic roles (Gomez, 2004a). This, in turn, has required from us to look very closely into the WordNet ontology for nouns and to modify it (Gomez, 2004b), when needed, based on the feedback obtained by running our verb predicate definitions through our semantic interpretation algorithm. The definition of the selectional restrictions has been a difficult task because we have to stay away from over-specifying or under-specifying them in handling figuratives senses of the verbs. In the following Framenet annotation:

"(A policeman) <DEFORMER> attempted to open the bathroom door with a crowbar , DENTING (the wood) <UNDERGOER> . (The experience) <DEFORMER> would not DENT (her confidence) <UNDERGOER> or make her fearful about driving alone, she said."

There is no distinction between the two senses of "dent" in those sentences, and, as a consequence, the different meanings of the roles DEFORMER and UNDERGOER. In our approach, the recognition of figurative senses of verbs is essential, and this has made much harder to define the selectional restrictions for the semantic roles of the predicates.

Another major difference between our work and Framenet is that we have approached the definition of verb predicates in a systematic way. As of this writing, we have defined over 3000 verb predicates and mapped about 98% of WordNet verb classes into verb predicates. The WordNet verb classes remaining contain only one or two verb forms. However, to our knowledge the verbs annotated in Framenet do not seem to follow any plan. Moreover, not all senses of the Framenet verbs seem to be annotated, and in many cases the annotators have selected verbs with very low polysemy, or just monosemous, becoming unclear which criteria the annotators have adopted in annotating some senses of a verb, and not others. In our case, we have approached the construction of our verb predicates classes by drawing on Fell-

baum's Wordnet verb classes (Fellbaum, 1998), which have been of an immense help to us. However, Fellbaum's concept of *troponymy* in many cases is too inclusive and as a result many verbs grouped under a Wordnet verb class do not share the same semantic roles, which is the major criterion used by us in building our verb predicate hierarchies. As a consequence, our ontology of verb predicates differs from Fellbaum's Wordnet verb classes since we have had to regroup them differently. Our verb predicates differ also from Levin (Levin, 1993) classes, which are based on syntactic criteria.

Another difference between our work and Framenet is that, in our view, there is an atomization of semantic roles in Framenet, e.g., "message," "messenger," "denter," "undergoer," "ingestibles," "ingestor," etc. In our approach, we have stayed away from reductionist approaches to verb semantics (see Jackendoff'localist hypothesis (Jackendoff, 1990), Dowty-Vendler's aspectual hypothesis (Vendler, 1967), Dowty (Dowty, 1979), or to a small set of primitives (Schank, 1975)) and also from approaches that lead to role propagation. An atomization of roles leads to an atomization of inferences that would be very hard to handle. Our approach does not lead to a propagation of roles since their number and nature depend on the generic predicate and its subpredicates.

However, we admit that, in some cases, it makes sense to recognize specific roles, or subroles as discussed in section 4.

8 Conclusions

We have described the semantic annotation of sentences with semantic roles carried out by a semantic interpreter. A corpus of 2000 interpreted sentences has been produced using the semantic interpreter. The main purpose in building this small corpus of 2000 sentences has been the validation of the verb predicate definitions. These definitions allow the semantic interpreter to determine verb meaning, or verb predicate, and its associated semantic roles. We have briefly described the verb predicates on which the semantic interpreter is based, the role of selectional restrictions, and its relation to the noun ontology. We have also provided a comparison to other approaches to annotation of semantic roles.

References

D. R. Dowty. 1979. *Word Meaning and Montague Grammar*. Dordrecht: Reidel, Dordrecht.

C. Fellbaum. 1998. A semantic network of english verbs. In C. Fellbaum, editor, *WordNet: An Electronic Lexical Database*, pages 69–104. MIT Press, 1998, Cambridge, Mass.

C. J. Fillmore, Johnson, C. R., and M.R.L. Petruck. 2003. Background to framenet. *International Journal of Lexicography*, 16:235–250.

F. Gomez. 1998. Linking wordnet verb classes to semantic interpretation. In *COLING-ACL Workshop on Usage of WordNet in NLP Systems*, pages 58–64, U. de Montreal.

F. Gomez. 2001. An algorithm for aspects of semantic interpretation using an enhanced wordnet. In *Proccedings of the 2nd North American Meeting of the North American Association for Computational Linguistics, NAACL-2001*, pages 87–94.

F. Gomez. 2004a. Building verb predicates: A computational view. In *Proccedings of the 42nd Meeting of the Association for Computational Linguistics, ACL-04*, pages 351–358, Barcelona, Spain.

F. Gomez. 2004b. Grounding the ontology on the semantic interpretation algorithm. In *Proceedings of the Second International WordNet Conference*, pages 124–129, Masaryk University, Brno.

R. Jackendoff. 1990. *Semantic Structures*. MIT Press, Cambridge, Mass.

P. Kingsbury, M. Palmer, and M. Marcus. 2002. Adding semantic annotation to the penn treebank. In *Proc. of the Human Language Technology Conference*, pages –, San Diego, CA.

B. Levin. 1993. *English Verb Classes and Alternations: A Preliminary Investigation*. University of Chicago Press, Chicago.

George Miller. 1998. Nouns in wordnet. In C. Fellbaum, editor, *WordNet: An Electronic Lexical Database*, pages 23–46. MIT Press, 1998, Cambridge, Mass.

R. Schank. 1975. *Conceptual Information Processing*. North Holland, Amsterdam.

Z. Vendler. 1967. *Linguistics in Philosophy*. Cornell University Press, Ithaca, New York.