The Deep Lexical Semantics of Event Words

Jerry R. Hobbs and Niloofar Montazeri
Information Sciences Institute
University of Southern California
Marina del Rey, California

1 Introduction

From the sentence

>   Russia is blocking oil from entering Ukraine.

we would like to be able to conclude

>   Oil can not be delivered to Ukraine.

But doing this requires fairly complex inference, because the words "block", "enter", "can", "not" and "deliver" carve up the world in different ways.[1]

Words describe the world, so if we are going to draw the appropriate inferences in understanding a text, we must have underlying theories of aspects of the world and we must have axioms that link these to words. This includes domain-dependent knowledge, of course, but 70-80% of the words in most texts, even technical texts, are words in ordinary English used with their ordinary meanings, like "enter" and "deliver". For example, so far in this paragraph, only the words "theories" and "axioms" and possibly "domain-dependent" have been domain-dependent.

Domain-independent words have such wide utility because their basic meanings tend to be very abstract, and they acquire more specific meanings in combination with their context. Therefore, the underlying theories required for explicating the meanings of these words are going to be very abstract.

For example, a core theory of scales will provide axioms involving predicates such as `scale`, `lessThan`, `subscale`, `top`, `bottom`, and `at`. These are abstract notions that apply to partial orderings as diverse as heights, money, and degrees of happiness. Then, at the "lexical periphery" we will be able to define the rather complex word "range" by the following axiom:[2]

```
(forall (x y z)
   (iff (range x y z)
        (exist (s s1 u1 u2)
           (and (scale s)(subscale s1 s)(bottom y s1)(top z s1)
                (member u1 x)(at u1 y)(member u2 x)(at u2 z)
                (forall (u)
                   (if (member u x)
                       (exist (v) (and (in v s1)(at u v))))))))))
```

[2]In this paper we use a subset of Common Logic (`http://common-logic.org/`)for the syntax of our notation.

That is, `x` ranges from `y` to `z` if and only if there is a scale `s` with a subscale `s1` whose bottom is `y` and whose top is `z`, such that some member `u1` of `x` is at `y`, some member `u2` of `x` is at `z`, and every member `u` of `x` is at some point `v` in `s1`.

Many things can be conceptualized as scales, and when this is done, a large vocabulary, including the word "range", becomes available. For example, we can now use and interpret "range" in the sentences

>The grades on the midterm ranged from 33 to 96.
>The timber wolf ranges from New Mexico to Alberta.
>Pat's behavior ranges from barely tolerable to deeply hostile.

by instantiating `scale` and `at` in different ways. From the axiom and the first sentence, we should be able to answer the questions

>Did someone get a 33 on the test? Yes.
>Did someone get a 22 on the test? No.
>Did someone get a 44 on the test? Maybe.

Similar questions could be answered for the other two sentences.

It would be good if we could learn relevant lexical and world knowledge automatically, and there has been some excellent work in this area (e.g., Pantel and Lin, 2002). For example, we can automatically learn the correlation between "married" and "divorced", and maybe we can even learn automatically the corresponding predicate-argument structures and which way the implication goes and with what temporal constraints. But this is a very simple relation to axiomatize in comparison to the "range" axiom. The kinds of knowledge we need are in general much more complex than automatic methods can give us. Moreover, automatic methods do not always yield very reliable results. The word "married" is highly correlated with "divorced" but it is also highly correlated with "murdered".

We are engaged in an enterprise we call "deep lexical semantics", in which we develop various core theories of fundamental commonsense phenomena and define English word senses by means of axioms using predicates explicated in these theories. Among the core theories is a theory of the structure of events, which is the focus of this paper.

If we construct the core theories and the linking axioms manually, we can achieve the desired complexity and reliability. However, it would not be feasible to axiomatize the meanings of 100,000 words manually. But it *is* feasible to axiomatize the meanings of several thousand words manually, and if the words are very common, this would result in a very valuable resource for natural language understanding.

We use textual entailment pairs like the "delivered" example above to test out subsets of related axioms. This process enforces a uniformity in the way axioms are constructed, and also exposes missing inferences in the core theories, as we discuss later in this chapter.

This chapter describes an effort in which a set of very common words somehow related to events and their structure are being linked with underlying core theories. Section 2 describes previous work in identifying a "core WordNet" and subsequent efforts to examine and classify the words in various ways. This led to the identification of 446 common words with senses that are primarily focused on events, viewed abstractly. In Section 3 we describe two aspects of the framework we are working in—the logical form we use, and abductive interpretation and defeasibility. In Section 4 we describe several of the core theories that are crucial in characterizing event words, including composite entities, scales, change, and causality. In Section 5 we describe the methodology we use for constructing axioms, deriving from WordNet and FrameNet senses a smaller set of abstract, general "supersenses", encoding axioms for these, and testing them on textual entailment pairs; we give as examples the analyses of several common words. In Section 6 we look at a specific example to illustrate both the power of the method for textual entailment and the holes in the knowledge base that it exposes. In Section 7 we address the problem of holes more systematically, specifically asking, for example, what kinds of "pairwise interactions" are possible for core theory predicates like `change` and `cause`.

2 Identifying the Core Event Words

WordNet (Miller, 1995) contains tens of thousands of synsets referring to highly specific animals, plants, chemical compounds, French mathematicians, and so on. Most of these are rarely relevant to any particular natural language understanding application. To focus on the more central words in English, the Princeton WordNet group has compiled a CoreWordNet[3], consisting of 4,979 synsets that express frequent and salient concepts. These were selected as follows: First, a list with the most frequent strings from the British National Corpus was automatically compiled and all WordNet synsets for these strings were pulled out. Second, two raters determined which of the senses of these strings expressed "salient" concepts (Boyd-Graber et al., 2006). Only nouns, verbs and adjectives were identified in this effort, but subsequently 322 adverbs were added to the list.

We classified these word senses manually into sixteen broad categories, listed here with rough descriptions and lists of sample words in the categories. Word senses are not indicated but should be obvious from the category.

Composite Entities: the structure and function of things made of other things: perfect, empty, relative, secondary, similar, odd, . . .

Scales: partial orderings and their fine-grained structure: step, degree, level, intensify, high, major, considerable, . . .

Events: concepts involving change and causality: constraint, secure, generate, fix, power, development, . . .

Space: spatial properties and relations: inside, top, list, direction, turn, enlarge, long, . . .

Time: temporal properties and relations: year, day, summer, recent, old, early, present, then, often, . . .

Cognition: concepts involving mental and emotional states: imagination, horror, rely, remind, matter, estimate, idea, . . .

Communication: concepts involving people communicating with each other: journal, poetry, announcement, gesture, charter, . . .

Persons: concepts involving persons and their relationships and activities: leisure, childhood, glance, cousin, jump, . . .

Microsocial: social phenomena other than communication that would be present in any society regardless of their level of technology: virtue, separate, friendly, married, company, name, . . .

Bio: living things other than humans: oak, shell, lion, eagle, shark, snail, fur, flock, . . .

Geo: geographical, geological and meteorological concepts: storm, moon, pole, world, peak, site, sea, island, . . .

Material World: other aspects of the natural world: smoke, stick, carbon, blue, burn, dry, tough, . . .

Artifacts: physical objects built by humans to fulfill some function: bell, button, van, shelf, machine, film, floor, glass, chair, . . .

Food: concepts involving things that are eaten or drunk: cheese, potato, milk, bread, cake, meat, beer, bake, spoil, . . .

Macrosocial: concepts that depend on a large-scale technological society: architecture, airport, headquarters, prosecution, . . .

Economic: having to do with money and trade: import, money, policy, poverty, profit, venture, owe, . . .

These categories of course have fuzzy boundaries and overlaps, but their purpose is only for grouping together concepts that need to be axiomatized together for coherent theories.

Each of these categories was then given a finer-grained structure. The internal structure of the category of event words is given below, with descriptions and examples of each subcategory.

---

[3]CoreWordNet is downloadable from http://wordnet.cs.princeton.edu/downloads.html.

- **State:** Having to do with an entity being in some state or not: have, remain, lack, still, . . .

- **Change**: involving a change of state:

  - Abstractly: incident, happen
  - A change of real or metaphorical position: enter, return, take, leave, rise, . . .
  - A change in real or metaphorical size or quantity: increase, fall, . . .
  - A change in property: change, become, transition, . . .
  - A change in existence: develop, revival, decay, break, . . .
  - A change in real or metaphorical possession: accumulation, fill, recovery, loss, give, . . .
  - The beginning of a change: source, start, origin, . . .
  - The end of a change: end, target, conclusion, stop, . . .
  - Things happening in the middle of a change: path, variation, repetition, [take a] break, . . .
  - Participant in a change: participant, player, . . .

- **Cause:** having to do with something causing or not causing a change of state:

  - In general: effect, result, make, prevent, so, thereby, . . .
  - Causes acting as a barrier: restriction, limit, restraint, . . .
  - An absence of causes or barriers: chance, accident, freely, . . .
  - Causing a change in position: put, pull, deliver, load, . . .
  - Causing a change in existence: develop, create, establish, . . .
  - Causing a change in real or metaphorical possession: obtain, deprive, . . .

- **Instrumentality:** involving causal factors intermediate between the primary cause and the primary effect: way, method, ability, influence, preparation, help, somehow, . . .

- **Process:** A complex of causally related changes of state:

  - The process as a whole: process, routine, work, operational, . . .
  - The beginning of the process: undertake, activate, ready, . . .
  - The end of the process: settlement, close, finish, . . .
  - Things that happen in the middle of a process: trend, steady, postpone, drift, . . .

- **Opposition:**

  - Involving factors acting against some causal flow: opposition, conflict, delay, block, bar, . . .
  - Involving resistance to opposition: resist, endure, . . .

- **Force:** Involving forces acting causally with greater or lesser intensity: power, strong, difficulty, throw, press, . . .

- **Functionality:** A notion of functionality with respect to some human agent's goals is superimposed on the causal structure; some outcomes are good and some are bad:

  - Relative to achieving a goal: use, success, improve, safe, . . .
  - Relative to failing to achieve a goal: failure, blow, disaster, critical, . . .
  - Relative to countering the failure to achieve a goal: survivor, escape, fix, reform, . . .

As with the broad categories, these subcategories are intended to group together words that need to be defined or characterized together if a coherent theory is to result.

3 Framework

Logical Notation: We use a logical notation in which states and events (eventualities) are reified. Specifically, if the expression `(p x)` says that `p` is true of `x`, then `(p' e x)` says that `e` is the eventuality of `p` being true of `x`. Eventuality `e` may exist in the real world (`Rexist`), in which case `(p x)` holds, or it may only exist in some modal context, in which case that is expressed simply as another property of the possible individual `e`.

The logical form of a sentence is a flat conjunction of existentially quantified postive literals, with about one literal per morpheme. (For example, logical words like "not" and "or" are treated as expressing predications about possible eventualities.) We have developed software[4] to translate Penn TreeBank-style trees (as well as other syntactic formalisms) into this notation. The underlying core theories are expressed as axioms in this notation (Hobbs, 1985).

As axiomatized, eventualities are isomorphic to predications, and just as predications have arguments, eventualities have participants. The expression `(arg x e)` says that entity `x` is a participant in or argument of eventuality `e`. We can define a predicate `relatedTo` that holds between two entities `x` and `y` when they are participants in the same eventuality, or equivalently, when they are arguments of the same predication.

We find that reifying states and events as eventualities and treating them as first-class individuals is preferable to employing the event calculus (Gruninger and Menzel, 2010; Mueller, 2006) which makes a sharp distinction between the two, because language makes no distinction in where they can appear and we can give them a uniform treatment.

Abduction: The interpretation of a text is taken to be the lowest-cost abductive proof of the logical form of the text, given the knowledge base. That is, to interpret a text we prove the logical form, allowing for assumptions at cost, and pick the lowest-cost proof. Factors involved in computing costs include, besides the number of assumptions, the salience of axioms, the plausibility of axioms expressing defeasible knowledge, and consiliance or the degree to which the pervasive implicit redundancy of natural language texts is exploited. We have demonstrated that many interpretation problems are solved as a by-product of finding the lowest-cost proof. This method has been implemented in an abductive theorem-prover called Mini-Tacitus[5] that has been used in a number of applications (Hobbs et al., 1993; Mulkar et al., 2007), and is used in the textual entailment problems described here.

Most commonsense knowledge is defeasible, i.e., it can be defeated. This is represented in our framework by having a unique "et cetera" proposition in the antecedent of Horn clauses that cannot be proved but can be assumed at a cost corresponding to the likeliehood that the conclusion is true. For example, the axiom

```
(forall (x) (if (and (bird x)(etc-i x))(fly x)))
```

would say that if `x` is a bird and other unspecified conditions hold, `(etc-i)`, then `x` flies. No other axioms enable proving `(etc-i x)`, but it can be assumed, and hence participate in the lowest cost proof. The index `i` is unique to this axiom. In this paper rather than invent new indices for each axiom, we will use the abbreviation `(etc)` to indicate the defeasibility of the rule. (This approach to defeasibility is similar to circumscription (McCarthy, 1980).)

4 Some Core Theories

The enterprise is to link words with core theories. In Section 2 gave an indication of the words involved in the effort, and a high-level analysis of the concepts needed for defining or characterizing them formally. This section sketches some of the principal core theories, including concepts used in Section 5.[6]

---

[4]`http://www.rutumulkar.com/download/NL-Pipeline/NL-Pipeline.php`.

[5]`http://rutumulkar.com/download/TACITUS/tacitus.php`.

[6]Descriptions of all the core theories, with axioms, can be found at `http://www.isi.edu/ hobbs/csk.html`.

Currently, there are sixteen theories defining or characterizing 230 predicates with 380 axioms. The theories differ from other commonsense knowledge bases, such as Cyc (Guha and Lenat, 1990) or SUMO (Niles and Pease, 2001), primarily in the abstract character and linguistic motivation of the knowledge.

Set Theory: This is axiomatized in a standard fashion, and provides predicates like *setdiff* and *deleteElt*, the latter expressing a relation between a set and the set resulting from deleting an element from it.

Composite Entities: This is a very general theory of things made of other things, one of the most basic notions one can imagine. A composite entity is characterized by a set of components, a set of properties of these components, and a set of relations among the components and between the components and the whole. With this theory we can talk about the structure of an entity by explicating its components and their relations, and we can talk about the environment of an entity by viewing the environment as composite and having the entity among its components. The predicate *partOf* is a very broad notion covering among other relations the *componentOf* relation. We also introduce in this theory the figure-ground relation at which places an external entity "at" some component in a composite entity.

Scales: This theory was mentioned in the introduction. In addition to defining the basic vocabulary for talking about partial orderings, we also explicate monotone-increasing scale-to-scale functions ("the more . . . the more . . ."), the construction of composite scales, the characterization of qualitatively high and low regions of a scale (related to distributions and functionality), and constraints on vague scales based on associated subsets (e.g., if Pat has all the skills Chris has and then some, Pat is more skilled than Chris, even though such judgments in general are often indeterminate).

Change of State: The two core theories most relevant to this chapter are the theory of change of state and the theory of causality. The predication (change' e e1 e2) says that e is a change of state whose initial state is e1 and whose final state is e2. The chief properties of change are that there is some entity whose state is undergoing change, that change is defeasibly transitive, that e1 and e2 cannot be the same unless there has been an intermediate state that is different, and that change is consistent with the before relation from our core theory of time. Since many lexical items focus only on the initial or the final state of a change, we introduce for convenience the predications (changeFrom' e e1) and (changeTo' e e2), defined in terms of change.

Cause: The chief distinction in our core theory of causality is between the notions of causalComplex and cause. A causal complex includes all the states and events that have to happen or hold in order for the effect to happen. A cause is that contextually relevant element of the causal complex that is somehow central to the effect, whether because it is an action the agent performs, because it is not normally true, or for some other reason. Most of our knowledge about causality is expressed in terms of the predicate cause, rather than in terms of causal complexes, because we rarely if ever know the complete causal complex. Typically planning, explanation, and the interpretation of texts (though not diagnosis) involves reasoning about cause. Among the principal properties of cause are that it is defeasibly transitive, that events defeasibly have causes, and that cause is consistent with before. In addition, in this theory we define such concepts as enable, prevent, help, and obstruct. There are also treatments of attempts, success, failure, ability, and difficulty.

Events: This theory is about how changes of state and causality compose into more complex events, processes and scenarios. It includes definitions of conditional, iterative, cyclic, and periodic events, and is linked with several well-developed ontologies for event structure, e.g., PSL (Bock and Grunninger, 2005).

Time: We also have a core theory of time, and the times of states and events can be represented as temporal properties of the reified eventualities. The theory of time has an essential function in axioms for words explicitly referencing time, such as "schedule" and "delay". But for most of the words we are explicating in this effort, we base our approach to the dynamic aspects of the world on the cognitively more basic theory of change of state. For example, the word "enter" is axiomatized as a change of state from being outside to being inside, and the fact that being outside comes *before* being inside follows from the axiom relating the predicates change and before.

5   Analyzing and Axiomatizing Word Senses

Our methodology consists of three steps.

1. Analyze the radial structure of a word's WordNet senses.

2. Write axioms for the most general senses.

3. Test the axioms on textual entailment pairs.

Our focus in this paper is on words involving the concepts of change of state and causality, or event words, such as "block", "delay", "deliver", "destroy", "enter", "escape", "give", "have", "hit", "manage", "provide", "remain", and "remove". For each word, we analyze the structure of its WordNet senses. Typically, there will be pairs that differ only in, for example, constraints on their arguments or in that one is inchoative and the other causative. This analysis generally leads to a radial structure indicating how one sense leads by increments, logically and perhaps chronologically, to another word sense (Lakoff, 1987). The analysis also leads us to posit "supersenses" that cover two or more WordNet senses. (Frequently, these supersenses correspond to senses in FrameNet (Baker et al., 2003) or VerbNet (Kipper et al., 2006), which tend to be coarser grained; sometimes the desired senses are in WordNet itself.)

"Enter": For example, for the verb "enter", three WordNet senses involve a change into a state:

V2: become a participant: "enter a race"
V4: play a part in: "this factor enters into your decision"
V9: set out on an enterprise: "enter a new career"

Call this supersense S1. Two other senses add a causal role to this (S2):

V5: make a record of: "enter the data"
V8: put or introduce into something: "enter a figure into a text"

Two more senses specialize supersense S1 by restricting the target state to be in a physical location (S1.1):

V1: come or go into: "he entered the room"
V6: come on stage: "enter from stage left"

One other sense specializes S1 by restricting the target state to be membership in a group (S1.2).

V3: register formally as a participant or member: "enter a club"

Knowing this radial structure of the senses helps enforce uniformity in the construction of the axioms. If the senses are close, their axioms should be almost the same.

Figure 1 shows the radial structure of the senses for the word "enter", together with the axioms that characterize each sense. A link between two word senses means an incremental change in the axiom for one gives the axiom for the other. For example, the axiom for `enter-S2` says that if `x1` enters `x2` in `x3`, then `x1` causes a change to the eventuality `i1` in which `x2` is in `x3`; and the expanded axiom for `enter-S1.1` states that if `x1` enters `x2`, then there is a change to a state `e1` in which `x1` is in `x2`. So `enter-S2` and `enter-S1.1` are closely related and thus linked together.

Abstraction is a special incremental change where one sense S1.1 specializes another sense S1 either by adding more predicates to or specializing some of the predicates in S1's axiom. We represent abstractions via arrows pointing from the subsenses to the supersenses. In Figure 1, `enter-S1.1` and `enter-S1.2` both specialize `enter-S1`. The predicate `enter-S1.1` adds an extra predicate describing `e1` as an `in` eventuality and `enter-S1.2` specializes `e1` to membership in `x2`, where `x2` is a group.

"Have": In WordNet the verb "have" has 19 senses. But they can be grouped into three broad supersenses. In its first supersense, X has Y means that X is in some relation to Y. The WordNet senses this covers are as follows:
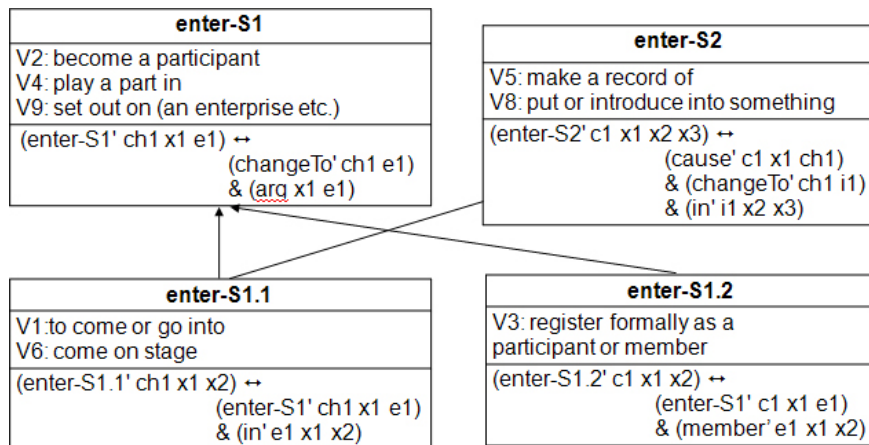
Figure 1: Senses of and axioms for the verb "enter"

    V1. a broad sense, including have a son, having a condition hold and having a college degree
    V2. having a feature or property, i.e., the property holding of the entity
    V3. a sentient being having a feeling or internal property
    V4. a person owning a possession
    V7. have a person related in some way: "have an assistant"
    V9. have left: "have three more chapters to write"
    V12. have a disease: "have influenza"
    V17. have a score in a game: "have three touchdowns"

The supersense can be characterized by the axiom

```
(forall (x y) (if (have-S1 x y)(relatedTo x y)))
```

(We use S suffixes for supersenses, W or V suffixes for WordNet senses, and F suffixes for FrameNet senses.)

The individual senses are then specializations of the supersense where more domain-specific predicates are explicated in more specialized domains. For example, sense 4 relates to the supersense as follows:

```
(forall (x y) (iff (have-W4 x y)(possess x y)))

(forall (x y) (if (have-W4 x y)(have-S1 x y)))
```

where the predicate `possess` would be explicated in a commonsense theory of economics, relating it to the priveleged use of the object. Similarly, `(have-W3 x y)` links with the supersense but has the restrictions that x is sentient and that the `relatedTo` property is the predicate-argument relation between the feeling and its subject.

The second supersense of "have" is "come to be in a relation to". This is our $changeTo$ predicate. Thus, the definition of this supersense is

```
(forall (x y)
   (iff (have-S2 x y)
        (exist (e) (and (changeTo e)(have-S1' e x y)))))
```

The WordNet senses this covers are as follows:

    V10. be confronted with: "we have a fine mess"
    V11. experience: "the stocks had a fast run-up"
    V14. receive something offered: "have this present"
    V15. come into possession of: "he had a gift from her"
    V16. undergo, e.g., an injury: "he had his arm broken in the fight"
    V18. have a baby

In these senses the new relation is initiated but the subject does not necessarily play a causal or agentive role. The particular change involved is specialized in the WordNet senses to a confronting, a receiving, a giving birth, and so on.

The third supersense of "have" is "cause to come to be in a relation to". The axiom defining this is

```
(forall (x y)
   (iff (have-S3 x y)
        (exist (e) (and (cause x e)(have-S2' e x y)))))
```

The WordNet senses this covers are

> V5. cause to move or be in a certain position or condition: "have your car ready"
> V6. consume: "have a cup of coffee"
> V8. organize: "have a party"
> V13. cause to do: "she had him see a doctor"
> V19. have sex with

In all these cases the subject initiates the change of state that occurs.

FrameNet has five simple transitive senses for "have". Their associated frames are

1. Have associated
2. Possession
3. Ingestion
4. Inclusion
5. Birth

The first sense corresponds to the first WordNet supersense:

```
(forall (x y) (iff (have-F1 x y)(have-S1 x y)))
```

The second sense is WordNet sense 4.

```
(forall (x y) (iff (have-F2 x y)(have-W4 x y)))
```

The third sense is WordNet sense 6. The fourth sense is the *partOf* relation introduced in Section 3. It is a specialization of WordNet sense 2.

```
(forall (x y) (iff (have-F4 x y)(partOf x y)))
```

```
(forall (x y) (if (have-F4 x y)(have-W2 x y)))
```

The fifth sense is WordNet sense 18.

By relating the senses in this way, an NLP system capable of inference can tap into both resources, for example, by accessing the WordNet hierarchy or the WordNet glosses expressed as logical axioms (Harabagiu and Moldovan, 2000), and by accessing the FrameNet frames, which are very close to axiomatic characterizations of abstract situations (cf. Ovchinnikova et al., 2011). In addition, it allows us to access the core theories explicating predicates like `relatedTo` and `partOf`.

"Remain:" There are four WordNet senses of the verb "remain":

> V1. Not change out of a state: "He remained calm."
> V2. Not change out of being at a location: "He remained at his post."
> V3. Entities in a set remaining after others are removed: "Three problems remain."
> V4. A condition remains in a location: "Some smoke remained after the fire was put out."

The first sense is the most general and subsumes the other three. We can characterize it by the axiom

```
(forall (x e)
   (if (remain-W1 x e)(and (arg x e)(not (changeFrom e)))))
```

By the properties of `changeFrom` it follows that `x` is in state `e`. In the second sense, the property `e` of `x` is being in a location.

```
(forall (x e)
   (iff (remain-W2 x e)
        (exist (y) (and (remain-W1 x e)(at' e y)))))
```

The fourth sense is a specialization of the second sense in which the entity x that remains is a state or condition.

The third sense is the most interesting to characterize. There is a process that removes elements from a set, and what remains is the set difference between the original and the set of elements that are removed. In this axiom x remains after process e.

```
(forall (x e)
   (iff (remain-W3 x e)
        (exist (y s1 s2 s3)
           (and (remove' e y s2 s1)(setdiff s3 s1 s2)
                (member x s3)))))
```

That is, x remains after e if and only if e is a removal event by some agent y of a subset s2 from s1, s3 is the set difference between s1 and s2, and x is a member of s3.

There are four FrameNet senses of "remain". The first is the same as WordNet sense 1. The second is the same as WordNet sense 3. The third and fourth are two specializations of WordNet sense 3, one in which the removal process is destructive and one in which it is not.

There are two nominalizations of the verb "remain"—"remainder" and "remains". All of their senses are related to WordNet sense 3. The first WordNet sense of "remainder" is the most general.

```
(forall (x y) (iff (remainder-W1 x e)(remain-W3 x e)))
```

That is, x is the remaineder after process e if and only if x remains after e. The other three senses result from specialization of the removal process to arithmetic division, arithmetic subtraction, and the purposeful cutting of a piece of cloth.

The supersenses capture the basic topology of the senses they subsume. The extra information that the subsenses convey are typically the types and properties of the arguments, such as being a place or a process, or qualities of the causing event, such as being sudden or forceful.

We are currently only constructing axioms for the most general or abstract senses or supersenses. In this way, although we are missing some of the implications of the more specialized senses, we are capturing the most basic topological structure in the meanings of the words. Moreover, the specialized senses usually tap into some specialized domain that needs to be axiomatized before the axioms for these senses can be written, e.g., ownership for have-W4.

In constructing the axioms in the event domain, we are very much informed by the long tradition of work on lexical decomposition in linguistics (e.g., Gruber, 1965; Jackendoff, 1972). Our work differs from this in that our decompositions are done as logical inferences and not as tree transformations as in the earliest linguistic work, they are not obligatory but only inferences that may or may not be part of the lowest-cost abductive proof, and the "primitives" into which we decompose the words are explicated in theories that enable reasoning about the concepts.

6   Textual Entailment

For each set of inferentially related words we construct textual entailment pairs, where the hypothesis (H) intuitively follows from text (T), and use these for testing and evaluation. The person writing the axioms does not know what the pairs are, and the person constructing the pairs does not know what the axioms look like.

The ideal test then is whether given a knowledge base K consisting of all the axioms, H cannot be proven from K alone, but H can be proven from the union of K and the best intepretation of T. This is often too stringent a condition, since H may contain irrelevant material that doesn't follow from T, so an alternative is to determine whether the lowest cost abductive proof of H given K plus T is substantially

lower than the lowest cost abductive proof of H given K alone, where "substantially lower" is defined by a threshold that can be trained (Ovchinnikova et al., 2011).

Here we work through an example to illustrate how textual entailment problems are handled in our framework. We assume in this example that lexical disambiguation has been done correctly. With more context, lexical disambiguation should fall out of the best interpretation, but it is unreasonable to expect that in these short examples. In practice we run the examples both with disambiguated and with nondisambiguated predicates. In this example we do not show the costs, although they are used by our system.

Consider the text-hypothesis pair we began with.

> T: Russia is blocking oil from entering Ukraine.
> H: Oil cannot be delivered to Ukraine.

What we notice in attempting to establish text-hypothesis relations like this after encoding the core theories and the axioms defining the words is that we get tantalizingly close to success, but not quite there, because of missing axioms. In Section 7 we discuss how this problem can be approached sysematically.

The relevant part of the logical form of the text is

```
(and (block-V3' b1 x1 e1)(enter-S2' e1 o1 u1))
```

That is, there is a blocking event `b1` in which Russia `x1` blocks eventuality `e1` from occurring, and `e1` is the eventuality of oil `o1` entering Ukraine `u1`. The `-V3` on `block` indicates that it is the third WordNet sense of the verb "block" and the `-S2` suffix on `enter` indicates that it is the second supersense of "enter".

The relevant part of the logical form of the hypothesis is

```
(and (not' n2 c2) (can-S1' c2 x2 d2) (deliver-S2' d2 x2 o2 u2))
```

That is, `n2` is the eventuality that `c2` is not the case, where `c2` is some `x2`'s being able to do `d2`, where `d2` is `x2`'s delivering oil `o2` to Ukraine `u2`. Note that we don't know yet that the oil and Ukraine in the two sentences are coreferential.

The axiom relating the third verb sense of "block" to the underlying core theories is

```
AX4: (forall (c1 x1 e1)
        (if (block-V3' c1 x1 e1)
          (exist (n1 p1)
            (and (cause' c1 x1 n1)(not' n1 p1)(possible' p1 e1)))))
```

This rule says that for `x1` to block some eventuality `e1` is for `x1` to cause `e1` not to be possible. (In this example, for expositional simplicity, we have allowed the eventuality `c1` of blocking be the same as the eventuality of causing, where properly they should be closely related but not identical.)

The other axioms needed in this example are

```
AX1: (forall (c1 e1)
        (if (and (possible' c1 e1)(etc))
          (exist (x1)(can-S1' c1 x1 e1))))

AX2: (forall (d1 x1 c1 r1 x2 x3)
        (if (and (cause' d1 x1 c1)(changeTo' c1 r1)(rel' r1 x2 x3)
          (deliver-S2' d1 x1 x2 x3))))

AX3: (forall (c1 x1 x2)
        (if (enter-S2' c1 x1 x2)
          (exist (i1)(and changeTo' c1 i1)(in' i1 x1 x2))))
```

AX1 says that defeasibly, if an eventuality `e1` is possible, then someone can do it. AX2 says that if `x1` causes a change to a situation `r1` in which `x2` in in some relation to `x3`, then in a very general sense (S2), `x1` has delivered `x2` to `x3`. AX3 says that if `c1` is the eventuality of `x1` entering `x2`, then `c1` is the change into a state `i1` in which `x1` is in `x2`.

Starting with the logical form of H as the initial interpretation and applying axioms AX1 and AX2, we get interpretation H1:

```
H1: (and (not' n2 c2) (possible' c2 d2) (cause' d2 x2 c1)
        (changeTo' c1 r1)(rel' r1 o2 u2))
```

At this point we are stuck in our effort to back-chain to T. An axiom is missing, namely, one that says that "in" is a relation between two entities.

```
AX5:  (forall (r x1 x2) (if (in' r1 x1 x2)(rel' r1 x1 x2)))
```

Using AX5, we can back-chain from H1 and derive interpretation H2:

```
H2: (and (not' n2 c2)(possible' c2 d2)(cause' d2 x2 c1)
        (changeTo' c1 r1)(in' r1 o2 u2))
```

We can then further back-chain with AX3 to interpretation H3:

```
H3: (and (not' n2 c2)(possible' c2 d2)(cause' d2 x2 c1)
        (enter-S2' c1  o2 u2))
```

Again, we need a missing axiom, AX6, to get closer to the logical form of T:

```
AX6:  (forall (p e1)
        (if (and (possible' p,e1)(etc))
            (exist (c x1) (and (possible' p c)(cause' c x1 e1)))))
```

That is, if something is possible, it is possible for something to cause it. Using this axiom, we can derive

```
H4: (and (not' n2 c2)(possible' c2 c1)(enter-S2' c1  o2 u2))
```

The final missing axiom, AX7, says that if x_1 causes eventuality c_2 not to occur, then c_2 doesn't occur.

```
AX7:  (forall (n x1 n1 c2)
        (if (and (cause' n x1 n1)(not' n1 c2))( not' n c2)))
```

Using this we derive interpretation H5.

```
H5: (and (cause' n2 x3 n)(not' n c2)(possible' c2 c1)(enter-S2' c1 o2 u2))
```

We can now apply the rule for "block", identifying `b1` and `n2`, `x1` and `x3`, `e1` and `c1`, `o1` and `o2`, and `u1` and `u2`, yielding H6 and establishing the entailment relation between H and T.

```
H6: (and (block-V3' n2 x3 c1)(enter-S2' c1 o2 u2))
```

It may seem at first blush that any new text-hypothesis pair will reveal new axioms that must be encoded, and that therefore it is hopeless ever to achieve completeness in the theories. But a closer examination reveals that the missing axioms all involve relations among the most fundamental predicates, like `cause`, `change`, `not`, and `possible`. These are axioms that should be a part of the core theories of change and causality. They are not a random collection of facts, any one of which may turn out to be necessary for any given example. Rather we can investigate the possibilities systematically. That investigation is what we describe in the following section.

7 Elaborating the Core Theories: Relations among Fundamental Predicates

For completeness in the core theories, we need to look at pairs of fundamental predicates and ask what relations hold between them, what their composition yields, and for each such axiom whether it is defeasible or indefeasible. The predicates we consider are `possible`, `Rexist`, `not`, `cause`, `changeFrom`, and `changeTo`.

The first type of axiom formulates the relationship between two predicates. For example, the rule relating `cause` and `Rexist` is

```
(forall (x e) (if (cause x e)(Rexist e)))
```

That is, if something is caused, then it actually occurs. Other rules of this type are as follows:

```
(forall (x e) (if (Rexist e)(possible e)))

(forall (e) (if (and (Rexist e)(etc))(exist (x)(cause x e))))

(forall (e2)
   (if (changeTo e2)
       (exist (e1)(and (changeFrom e1)(not' e1 e2)))))

(forall (e1)
   (if (changeFrom e1)
       (exist (e2)(and (changeTo e2)(not' e2 e1)))))

(forall (e) (if (changeTo e)(Rexist e)))

(forall (e) (if (changeFrom e)(not e)))

(forall (e) (if (and (Rexist e)(etc))(changeTo e)))
```

That is, if something occurs, it is possible and, defeasibly, something causes it. If there is a change to some state obtaining, then there is a change from its not obtaining, and vice versa. If there is a change to something, then it obtains, and if there is a change from something, then it no longer obtains. If some state obtains, then defeasibly there was a change from something else to that state obtaining.

The second type of axiom involves the composition of predicates, and gives us rules of the form

```
(forall (e1 e2 x) (if (and (p' e1 e2)(q' e2 x)) (r' e1 x)))
```

That is, when p is applied to q, what relation r do we get?

Figure 2 shows the axioms encoding these compositions. The rows correspond to the (p' e1 e2)'s and the columns correspond to the (q' e2 x)'s, and the cell contains the consequents (r' e1 x). If the rule is defeasible, the cell indicates that by adding (etc) to the antecedent. The consequents in italics are derivable from other rules.

	(possible' e2 e3)	(Rexist' e2 e3)	(not' e2 e3)	(cause' e2 x2 e3)	(changeFrom' e2 e3)	(changeTo' e2 e3)
(possible' e1 e2)	(possible' e1 e3)	(possible' e1 e3)		*(possible' e1 e3)*		*(possible' e1 e3)*
(Rexist' e1 e2)	(possible' e1 e3)	(Rexist' e1 e3)	(not' e1 e3)	*(Rexist' e1 e3)*	*(not' e1 e3)*	*(Rexist' e1 e3)*
(not' e1 e2)	(not' e1 e3)	(not' e1 e3)	(Rexist' e1 e3)	(etc)->(not' e1 e3)	(Rexist' e1 e3)	(etc)->(not' e1 e3)
(cause' e1 x1 e2)	*(possible' e1 e3)*	(cause' e1 x1 e3) / *(Rexist' e1 e3)*	*(not' e1 e3)*	(cause' e1 x1 e3) / *(Rexist' e1 e3)*	*(not' e1 e3)* / *(changeFrom' e1 e3)*	*(cause' e1 x1 e3)* / *(changeTo' e1 e3)*
(changeFrom' e1 e2)		(changeFrom' e1 e3)	(changeTo' e1 e3) / *(Rexist' e1 e3)*	(etc)-> (changeFrom' e1 e3) / (etc)-> (not' e1 e3)	*(Rexist' e1 e3)*	(etc)->(not' e1 e3)
(changeTo' e1 e2)	*(possible' e1 e3)*	(changeTo' e1 e3) / *(Rexist' e1 e3)*	(changeFrom' e1 e3) / (not' e1 e3)	(etc)-> (changeTo' e1 e3) / (Rexist' e1 e3) / (cause' e1 x1 e3)	*(not' e1 e3)* / *(changeFrom' e1 e3)*	(Rexist' e1 e3) / *(changeTo' e1 e3)*

Figure 2: Axioms expressing compositions of fundamental predicates

For example, in the `possible-possible` cell, the rule says that if it is possible that something is possible, then it is possible. To take a more complex example, the `changeFrom-cause` cell says that if there is a change from some entity causing (or maintaining) a state, then defeasibly there will be a change from that state. So if a glass is released, it will fall.

We have also looked at axioms whose pattern is the converse of those in Figure 2. For example, if something does not hold, then it was not caused.

8 Summary

We understand language so well because we know so much, and our computer programs will only approach what we might call "understanding" when they have access to very large knowledge bases. Much of this knowledge will be of a technical nature and can perhaps be acquired automatically by statistical methods or from learning by reading. But the bulk of the inferences required for understanding natural language discourse involve very basic abstract categories. In the work described here, we have identified the words related to events and their structure, which because of their frequency are most demanding of explication in terms of the inferences they trigger. We have constructed abstract core theories of the principal domains that need to be elaborated in order to express these inferences in a coherent fashion. We presented a methodology for defining or characterizing the meanings of words in terms of the core theories, of evaluating the axioms using textul entailment, and of elaborating the knowledge base by identifying and filling lacunae. Doing this for several thousand of the most common words in English would produce a huge gain in the inferential power of our systems and would be an enterprise approximately equal in scope to the manual construction of other widely used resources such as WordNet and FrameNet. In combination with other knowledge resources, this work should take us a step closer to sophisticated, inference-based natural language understanding.

References

Baker, C., Fillmore, C., Cronin, B.: The Structure of the Framenet Database, International Journal of Lexicography, Volume 16.3: (2003) 281-296.

Bock, C, Gruninger, M.: PSL: A Semantic Domain for Flow Models, Software and Systems Modeling Journal, 4:2, (2005) 209 - 231.

Boyd-Graber, J., Fellbaum, C., Osherson, D., and Schapire, R.: Adding dense, weighted, connections to WordNet. In: Proceedings of the Third Global WordNet Meeting, Jeju Island, Korea, January 2006.

Gruber, Jeffrey C., 1965. *Studies in Lexical Relations*, unpublished Ph.D. dissertation, Massachusetts Institute of Technology, Cambridge, Massachusetts.

Gruninger, Michael, and Christopher Menzel, 2010. "The Process Specification Language (PSL) Theory and Applications", *AI Magazine*, Vol 24, No 3.

Guha, R., Lenat, D.: 1990. CYC: A Midterm Report. In: AI Magazine 11 (3), Fall 1990, pp. 33-59.

Harabagiu, S., Moldovan, D.: Enriching the WordNet Taxonomy with Contextual Knowledge Acquired from Text, in Natural Language Processing and Knowledge Representation: Language for Knowledge and Knowledge for Language, (Eds) S. Shapiro and L. Iwanska, AAAI/MIT Press, (2000) 301-334.

Hobbs, Jerry R. 1985. "Ontological Promiscuity." *Proceedings, 23rd Annual Meeting of the Association for Computational Linguistics*, pp. 61-69. Chicago, Illinois, July 1985.

Hobbs, Jerry R., 2008. "Deep Lexical Semantics", *Proceedings*, 9th International Conference on Intelligent Text Processing and Computational Linguistics (CICLing-2008), Haifa, Israel, February 2008.

Hobbs, Jerry R., Mark Stickel, Douglas Appelt, and Paul Martin, 1993. "Interpretation as Abduction", *Artificial Intelligence*, Vol. 63, Nos. 1-2, pp. 69-142.

Jackendoff, Ray S. 1972. *Semantic interpretation in generative grammar*. Cambridge, MA: The MIT Press.

Lakoff, George, 1987. *Women, Fire, and Dangerous Things: What Categories Reveal About the Mind*, University of Chicago Press, Chicago.

McCarthy, John, 1980. "Circumscription: A Form of Nonmonotonic Reasoning", *Artificial Intelligence*, 13: 27-39.

Miller, G.: WordNet: a lexical database for English. In: Communications of the ACM 38 (11), November 1995, pp. 39-41.

Mueller, Erik T., 2006. *Commonsense Reasoning*, Morgan Kaufmann Publishers, Inc., San Mateo, California.

R. Mulkar, J.R. Hobbs, and E. Hovy, 2007. "Learning from Reading Syntactically Complex Biology Texts", *Proceedings of the AAAI Spring Symposium Commonsense'07*. Stanford University, CA, 2007.

Niles, I., Pease, A.: Toward a Standard Upper Ontology, in Proceedings of the 2nd International Conference on Formal Ontology in Information Systems (FOIS-2001), (Eds) C. Welty and B. Smith, Ogunquit, Maine, October 2001.

E. Ovchinnikova, N. Montazeri, T. Alexandrov, J. Hobbs, M. McCord, and R. Mulkar-Mehta, 2011. "Abductive Reasoning with a Large Knowledge Base for Discourse Processing", *Proceedings of the 9th International Conference on Computational Semantics*, Oxford, United Kingdom.

Pantel, P., Lin, D.: 2002. Discovering Word Senses from Text. In Proceedings of ACM Conference on Knowledge Discovery and Data Mining (KDD-02). (2002) 613-619. Edmonton, Canada.