Reconciling fine-grained lexical knowledge and coarse-grained ontologies in the representation of near-synonyms

Philip Edmonds
Sharp Laboratories of Europe
Oxford Science Park, Edmund Halley Road
Oxford OX4 4GB
England
phil@sharp.co.uk

Graeme Hirst
Department of Computer Science
University of Toronto
Toronto, Ontario
Canada M5S 3G4
gh@cs.toronto.edu

Introduction

A machine translation system must be able to adequately cope with near-synonymy for there are often many slightly different translations available for any given source language word that can each significantly and differently affect the meaning or style of a translated text. Conventional models of lexical knowledge used in natural-language processing systems are inadequate for representing near-synonyms, because they are unable to represent fine-grained lexical knowledge. We will discuss a new model for representing fine-grained lexical knowledge whose basis is the idea of granularity of representation.

Near-synonymy

True synonyms, if they exist at all, are very rare. One would have to find two words that can be intersubstituted in all possible contexts (or expressions) without changing the meaning, however 'meaning' is defined. Philosophers such as Quine (1951) and Goodman (1952) argue that true synonymy is impossible because it is impossible to define, and so, perhaps unintentionally, dismiss all other forms of synonymy. Even if absolute synonymy were possible, pragmatic and empirical arguments show that it would be very rare. Cruse (1986, p. 270) says that "natural languages abhor absolute synonyms just as nature abhors a vacuum", because the meanings of words are constantly changing. More formally, Clark (1992) employs her Principle of Contrast, that "every two forms contrast in meaning", to show that language works to eliminate absolute synonyms. Either an absolute synonym would fall into disuse or it would take on a new nuance of meaning.

Near-synonyms, on the other hand, are pervasive. In fact, every so-called "dictionary of synonyms" actually contains only near-synonyms, a fact made obvious by those dictionaries, such as *Webster's New Dictionary of Synonyms* (Gove 1973) and Hayakawa's *Choose the Right Word* (Hayakawa 1994), that explicitly discriminate between words. Figure 1 shows a typical entry. Writers often turn to such resources when faced with a choice between near-synonyms, because choosing the wrong word can be imprecise or awkward, let alone convey an unwanted implication. For example, one would not refer to a someone else's *mistake* as a *blunder* unless one didn't mind the risk of offending them. That's because *blunder* is harsher than *mistake* and can imply ignorance and blameworthiness, whereas *mistake* implies mere misconception (Gove 1973; see figure 1).

Differences naturally come up when one talks about near-synonyms, for if two putative absolute synonyms aren't actually identical, then there must be something that makes them different. According to Saussure (1916, p. 114), difference is more important than similarity:

> In a given language, all the words which express neighbouring ideas help define one another's meaning. Each of a set of synonyms like *redouter* ('to dread'), *craindre* ('to fear'), *avoir peur* ('to be afraid') has its particular value only because they stand in contrast with one another. . . . No word has a value that can be identified independently of what else is in its vicinity.

Contrary to what one might expect—that the more similar two items are the easier it is to represent their differences, which is one interpretation of Markman and Gentner's work on similarity (Markman & Gentner 1993; Gentner & Markman 1994)—there is actually remarkable complexity in the differences between near-synonyms. Consider the full complexity of the usage note in figure 1. Not only do the near-synonyms differ in the expression of different concepts and ideas, such as misconception and blameworthiness, but also in the manner in which the concepts are conveyed (e.g., implied, suggested, expressed, connoted, and stressed), the frequency with which they are conveyed (e.g., commonly, sometimes, not always), and in the degree to which they are conveyed (e.g., in strength).

This example illustrates merely one broad type of variation, denotational variation. In general, near-synonyms can differ with respect to any aspect of their meaning. According to Cruse's (1986) analysis, the differentiation of near-synonyms can involve:

- collocational and syntactic variation,
- stylistic variation (including dialect and register),
- expressive variation (including emotive and attitudinal aspects), and
- denotational variation (in a broad sense, including propositional, fuzzy, and other peripheral aspects).

Continuing the work of DiMarco, Hirst, & Stede (1993) in analysing the types of differentiae used in synonym-

Error implies a straying from a proper course and suggests guilt as may lie in failure to take proper advantage of a guide . . .

Mistake implies misconception, misunderstanding, a wrong but not always blameworthy judgment, or inadvertence; it expresses less severe criticism than *error*.

Blunder is harsher than *mistake* or *error*; it commonly implies ignorance or stupidity, sometimes blameworthiness.

Slip carries a stronger implication of inadvertence or accident than *mistake*, and often, in addition, connotes triviality.

Lapse, though sometimes used interchangeably with *slip*, stresses forgetfulness, weakness, or inattention more than accident; thus, one says a *lapse* of memory or a *slip* of the pen, but not vice versa.

Faux pas is most frequently applied to a mistake in etiquette.

Bull, howler, and boner are rather informal terms applicable to blunders that typically have an amusing aspect.

Figure 1: An abridged entry from *Webster's New Dictionary of Synonyms* (Gove 1973).

Type of variation	Example
Collocational	*task* : *job*
Stylistic, formality	*pissed* : *drunk* : *inebriated*
Stylistic, force	*ruin* : *annihilate*
Expressed attitude	*skinny* : *thin* : *slim*
Emotive	*daddy* : *dad* : *father*
Abstract dimension	*seep* : *drip*
Emphasis	*enemy* : *foe*
Denotational, indirect	*error* : *blunder*
Denotational, fuzzy	*woods* : *forest*

Table 1: Examples of near-synonymic variation.

discrimination dictionaries, Edmonds (1999) gives a classification of near-synonymic variation into 35 subcategories of the above four broad categories. Table 1 gives several examples, which we will discuss briefly.

Collocational variation involves the words or concepts with which a word can be combined, possibly idiomatically, in a well-formed sentence. For example, *task* and *job* differ in their collocational patterns: *face a daunting job* sounds unnatural where *face a daunting task* does not. Stylistic variation involves differences in a small finite set of dimensions on which all words can be compared. Many stylistic dimensions have been proposed by Hovy (1988), Nirenburg & Defrise (1992), Stede (1993), and others. Table 1 illustrates two of the most common dimensions: *Inebriated* is formal while *pissed* is informal; *annihilate* is a more forceful way of saying *ruin*. Expressive variation can be used by a speaker to express his or her attitude towards a participant of the situation being spoken about. So *slim* can be a favourable description of someone's girth, whereas *skinny* can be pejorative, even though the denotation is the same. Notice that, unlike a stylistic distinction, which pertains strictly to the word itself, an expressed attitude explicitly involves a participant of the situation (e.g., the person that is *slim* or *skinny*). A word can also indirectly express the emotions of the speaker in a possibly finite set of emotive 'fields' (e.g., *daddy* expresses a stronger feeling of intimacy than *dad* or *father*).

There is a large range of phenomena involving denotational meaning, taken in a broad sense. DiMarco, Hirst, & Stede (1993) found that while some differentiae are easily expressed in terms of clear-cut abstract (or symbolic) features (such as `continuous/intermittent`: *Wine* {*seeped* | *dripped*} *from the barrel*), many are not. In fact, denotational variation involves mostly differences in the infinite range of concepts or ideas that words can convey (see figure 1), and, of course, the concepts are not simple features, but are full-fledged concepts that relate roles and aspects of a situation. (E.g., 'severity of criticism' is a complex concept that involves both a 'criticizer' and a 'criticizee', the one who made the error.)

Moreover, as we mentioned above, two words can differ in the manner in which they convey a concept. *Foe* and *enemy*, for instance, differ in emphasis, the former stressing antagonism and the latter active warfare rather than emotional reaction (Gove 1973). Other words convey meaning *indirectly* by mere suggestion or implication. There appears to be a continuum of indirectness from 'suggestion' to 'implication' to 'denotation'. Such indirect meanings are usually peripheral to the main meaning conveyed by an expression, and it is usually difficult to ascertain definitively whether or not they were even intended to be conveyed by the speaker. Differences in meaning can also be fuzzy, rather than clear-cut. The difference between *forest* and *woods* is a complex combination of size, primitiveness, proximity to civilization, and wildness (Room 1981).

Overall, differences between near-synonyms can be multi-dimensional, involving any or all of the types discussed above, and the dimensions need not be mutually independent. (E.g., *blunder* and *mistake* differ on several denotational dimensions, as well as on stylistic and attitudinal dimensions.)

Of course, near-synonyms can often be intersubstituted with no apparent change of effect on an utterance, but it is very difficult to specify just how the context influences the relevance of lexical differences. Lexicographers, for instance, resort to using 'frequency' terms such as *sometimes* and *usually* (see figure 1). So, while we can't yet solve the problems of context-dependence, we can actually work out the other representational problems in advance of having a proper theory of context.

Clearly, near-synonymy is a more interesting phenomenon than absolute synonymy because of the questions it raises about fine-grained lexical knowledge representation. But is near-synonymy a phenomenon in its own right warranting its own special account? One could, of course, argue that if there are no (absolute) synonyms, then even very similar words should be treated like any other pair of words, similar or not. However, in this paper we argue that

near-synonymy is indeed a separately characterizable phenomenon of word meaning. Our first evidence is that the conventional model of lexical knowledge as used in many computational systems is inadequate for representing near-synonyms.

Near-synonymy in lexical semantics

Current models of lexical knowledge used in computational systems, which are based on decompositional and relational theories of word meaning (Katz & Fodor 1963; Jackendoff 1990; Lyons 1977; Nirenburg & Defrise 1992; Lehrer & Kittay 1992b; Evens 1988; Cruse 1986), cannot account for the properties of near-synonyms. In outline, the now-conventional view is that the denotation of a lexical item is represented as a concept, or a structure of concepts (i.e., a word sense is linked to the concept that it lexicalizes), which are themselves organized into a taxonomy, often an ontology. The ontology is often language-independent, or at least language-neutral, so that it can be used in multilingual applications. Following Frege's (1892) or Tarski's (1944) truth-conditional semantics, the concept that a lexical item denotes can be thought of as a set of features that are individually necessary and collectively sufficient to define the concept. (Such a view greatly simplifies the word–concept link. In a text generation system, for instance, the features amount to the necessary applicability conditions of a word, i.e., they have to be present in the input in order to output the word.) So, except for polysemy and absolute synonymy, there is no logical difference between a word and a concept. Therefore, words that are nearly synonymous would have to be linked each to their own slightly different concepts.

While such models have been successful in computational systems, they are rarely pushed to represent near-synonyms—the work of Barnett, Mani, & Rich (1994) is a notable exception—but instead treat near-synonyms as absolute synonyms. Several researchers, including DiMarco, Hirst, & Stede (1993), Hirst (1995), and Edmonds (1999), discuss the inadequacies of such models with respect to near-synonymy. Briefly, the models have problems in the following three areas:

- The adequacy of coverage of phenomena related to near-synonymy.
- Engineering advantages, both in the design of an efficient and robust lexical choice process, and in the design of lexical entries for near-synonyms.
- The tractability of reasoning about concepts during natural language understanding and generation.

Hirst (1995) shows that a model in which every word, near-synonym or not, must be linked to its own concept entails an awkward taxonomic proliferation of language-specific concepts at the fringes of the ontology, thereby defeating the purpose of a language-independent ontology. Such a taxonomy could not efficiently represent the multi-dimensional nature of near-synonymic variation, nor could it account for fuzzy differences between near-synonyms.

A taxonomic hierarchy also emphasizes hyponymy, backgrounding all other relations, which appear to be more important in representing the multi-dimensional nature of fine-grained word meaning. It is not even clear that a group of synonyms *can* be structured by hyponymy.

And, since the model defines words in terms of only necessary and sufficient truth-conditions, it cannot account for indirect expressions of meaning and context-dependent meanings, which are clearly not necessary features of a word's meaning. It also cannot easily or tractably account for fuzzy differences or the full-fledged concepts required for representing denotational variation. (First-order logic would be required to represent such concepts, but reasoning about the concepts in lexical choice and other tasks would become intractable as more near-synonyms were represented.)

Nevertheless, to build a new model of lexical knowledge that takes into account the fine-grainedness of near-synonymy, the logical way forward is thus to start with the computationally-proven ontological model, and to modify or extend it to account for near-synonymy. Our new model will rely on a characterization of near-synonymy involving granularity of representation.

Near-synonymy and granularity of representation

Semanticists such as Ullmann (1962), Cruse (1986), and Lyons (1995) have attempted to define near-synonymy by focusing on 'propositional' meaning. Cruse's *cognitive synonyms*, for instance, are words that when intersubstituted in a sentence preserve its truth-conditions, but may change the expressive meaning, style, or register of the sentence, or may involve different idiosyncratic collocations (e.g., *violin* : *fiddle*).[1] His *plesionyms*, on the other hand, change the truth-conditions, but still yield semantically similar sentences (e.g., *misty* : *foggy*). Although these definitions are important for truth-conditional semantics, they are not very helpful for us, because plesionymy is left to handle all of the most interesting phenomena. Moreover, a rigorous definition of cognitive synonymy is difficult to come up with, because it relies on the notion of granularity, which we will discuss below.

Lexicographers have always treated synonymy as near-synonymy. They define synonymy in terms of likeness of meaning, disagreeing only in how broad the definition ought to be. For instance, Roget followed the vague principle of "the grouping of words according to ideas" (Chapman 1992, p. xiv). And since *Roget's Thesaurus* is structured hierarchically, words are ultimately grouped according to proximity of meaning: "the sequence of terms within a paragraph, far from being random, is determined by close, semantic relationships." (p. xiii). The lexicographers of *Webster's New Dictionary of Synonyms* define a synonym as "one of two or more words in the English language which have the same or very nearly the same *essential* meaning ... synonyms can be defined in the same terms up to a certain point" (Gove 1973, p. 24a). *Webster's Collegiate Thesaurus* uses a similar definition that involves the sharing of elementary mean-

[1] There is an old joke: What's the difference between a violin and a fiddle? No one minds if you spill beer on the fiddle.

ings, which are "discrete objective denotations uncolored by ... peripheral aspects such as connotations, implications, or quirks of idiomatic usage" (Kay 1988, p. 9a). Clearly, the main point of these definitions is that near-synonyms must have the same essential meaning, but may differ in peripheral or subordinate ideas. Cruse (1986, p. 267) actually refines this idea and suggests that synonyms, of all types, are words that are identical in "central semantic traits" and differ, if at all, in only "peripheral traits". But specifying formally just how much similarity of central traits and dissimilarity of peripheral traits is allowed can be a problem. That is, just what counts as a central trait and as a peripheral trait in defining a word? To discuss this further, we will now introduce the idea of granularity of representation.

The key to near-synonymy lies in the idea of *granularity of representation* of word meaning. By granularity we mean the level of detail used to describe or represent the meanings of words. A fine-grained representation can encode subtle distinctions, whereas a coarse-grained representation is crude and glosses over variation.

Granularity is independent of specificity, which is a property of concepts rather than representations of concepts. But there is a relationship between specificity and granularity, because only more-general concepts would take part in a coarse-grained representation of a concept. For example, a very general concept, say `Human`, can have, in a particular system, a very fine-grained representation, involving, say, a detailed description of the appearance of a human, references to related concepts such as `Eat` and `Procreate`, and information to distinguish the concept from other similar concepts such as `Animal`. Conversely, a very specific concept can have a very coarse-grained representation, which could involve only very general concepts; we could represent a `Tree` at a level of detail only enough to say that it is a physical object.

Near-synonyms can occur at any level of specificity, but crucially it is the granularity of the representations of their meanings that enables one to distinguish one near-synonym from another. So, any definition of synonymy that does not take granularity into account is insufficient. Take Cruse's cognitive synonymy, discussed above. On the one hand, at an absurdly coarse grain of representation, any two words could be considered cognitive synonyms (because every word denotes a 'thing'). And on the other, no two words could ever be cognitive synonyms, because, even at a fine grain, apparent cognitive synonyms may be further distinguishable by a still more fine-grained representation. For instance, *drunk* and *inebriated* may seem to differ in only non-propositional traits, but *inebriated* can imply an intoxication that results in exhilaration (Gove 1973), indicating that the words are in fact distinguishable at a finer grain. Thus, which pairs of words are cognitive synonyms depends on the granularity with which we represent their propositional meanings.

By taking granularity into account, we can come up with a much more useful definition of near-synonymy, because we can now characterize the difference between *essential* and *peripheral* aspects of meaning. Assuming that we can set an appropriate level of granularity, the essential meaning of a word is the portion of its meaning that is representable only above that level of granularity, and peripheral meanings are representable only below that level.

But what is the appropriate level of granularity, the dividing line, as it were, between coarse-grained and fine-grained representations? We could simply use our intuition, or rather, the intuitions of lexicographers, which are filtered by some amount of objectivity and experience.

Alternatively, from a concern for the representation of lexical knowledge in a multilingual application, we can view words as (obviously, language-specific) specializations of language-independent concepts. For example, DiMarco, Hirst, & Stede (1993) and Hirst (1995) presume a hierarchical organization of coarse-grained language-independent concepts. A set of near-synonyms is simply a set of words that all link to the same language-independent concept. So, in effect, near-synonyms share the same propositional meaning up to a point in granularity defined by language-dependence. Thus, if we are designing a lexical resource for use in multilingual applications, we have an operational definition of near-synonymy: if the same concept has several reasonable lexicalizations in different languages then it is a good candidate for being considered a language-independent concept, its various lexicalizations forming sets of near-synonyms in each language. Of course, in a unilingual application, we would have to come up with some other, perhaps domain-dependent, cut-off point in granularity (though the multilingual principle could also be used in such an application).

Granularity also explains why it is more difficult to represent near-synonyms in a lexicon. Near-synonyms are so close in meaning, sharing all essential coarse-grained aspects, that they differ, by definition, in only aspects representable at a fine grain. And these fine-grained representations of differences would tend to involve very specific concepts, which could involve complex structures of more general concepts. Such structures are more difficult to represent and to reason with. The matter is made only more complicated by there often being several interrelated near-synonyms with interrelated differences. On the other hand, words that are not near-synonyms—that are merely similar or not similar at all—could presumably be differentiated by concepts at a coarse-grained, and less complex, level of representation.

The above discussion has been leading to a new model of lexical knowledge representation in which near-synonymy is handled on a separate level of representation.

A model of fine-grained lexical knowledge

Our model is based on our contention that the meaning of a word, however it manifests itself in text or speech, *arises* out of a context-dependent combination of a basic inherent context-independent denotation and a set of explicit differences to its near-synonyms. (We don't rule out other elements in the combination, but these are the main two.) Thus, word meaning is not explicitly represented in the lexicon, but is created (or generated, as in a generative model of the

lexicon; see (Pustejovsky 1995)) when a word is used (i.e., during processing by a system). This theory preserves some aspects of the classical theories—the basic denotation can be modelled using an ontology—but the rest of a word's meaning relies on other nearby words and the context of usage (cf. Saussure). And since the relations of a word to its near-synonyms are so important to its meaning—possibly more important than other relations—near-synonyms form *clusters*.[2]

The theory is built on the following three premises, which follow from our observations about near-synonymy. First, the meaning of any word, at some level of granularity, must indeed have some inherent context-independent denotational aspect to it—if it did not, one would not be able to define or understand a word in isolation, as one in fact can. Second, nuances of meaning, while difficult or impossible to represent in positive, absolute, and context-independent terms, can be represented as differences, in Saussure's sense, between near-synonyms. That is, every nuance of meaning that a word might have can be thought of as a relation between the word and one or more of its near-synonyms. And third, differences must be described not as simple features or truth-conditions, but by structures that encode degrees of necessity, fuzziness, and relations to the context.

For example, the word *forest* denotes a geographical tract of trees at a coarse grain, but it is only in relation to *woods*, *copse*, and other near-synonyms that one can fully understand the significance of *forest* (i.e., that it is larger, more wild, etc.). Similarly, the word *mistake* denotes any sort of action that deviates from what is correct and also involves some notion of criticism, but only in relation to *error* and *blunder* does one see that the word can be used to criticize less severely than some alternative *error* words allow. None of these differences could be represented in absolute terms, because we would have to define some absolute notion of size, wildness, or severity, which seems implausible. So, at a fine grain, and only at a fine grain, we make explicit use of Saussure's notion of contrast in demarcating the meanings of near-synonyms. Hence, the theory holds that near-synonyms are explicitly related to each other not at a conceptual level but at a *subconceptual* level—outside of the (coarser-grained) ontology. In this way, a cluster of near-synonyms is not a mere list of synonyms; it has an internal structure that encodes fine-grained meaning as differences between lexical entries, and externally it is situated within a conceptual model (i.e., the ontology) on one side, and a linguistic model on the other side.

Specifically, we posit a model that has three levels of representation. Current computational theories suggest that two levels of representation, a *conceptual–semantic* level and a *syntactic–semantic* level, are necessary to account for various lexico-semantic phenomena in computational systems, including compositional phenomena such as paraphrasing (see, for instance, Stede's (1999) model). We believe a two-level semantics is also necessary, but to account for fine-grained meanings and near-synonymy, our theory postulates a third intermediate level (or a splitting of the conceptual–semantic level). The three levels are the following:

- A conceptual–semantic level.
- A subconceptual/stylistic–semantic level.
- A syntactic–semantic level.

At the top level, as in the current theories, we represent the coarse-grained essential denotational meaning of a word. At the intermediate level, we represent the fine-grained context-dependent differences in word meaning. At the bottom level, we represent how a word can be combined with other words both syntactically and collocationally.

So, if we were to take the conventional ontological model as a starting point, we would cut off the ontology at a coarse grain and cluster near-synonyms under a shared concept rather than linking each word to a separate concept. The resulting model is a *clustered model of lexical knowledge.* On the conceptual–semantic level, a cluster has a *core denotation* that represents the essential shared denotational meaning of its near-synonyms. On the subconceptual/stylistic–semantic level, we represent semantic, stylistic, and expressive *distinctions* between the near-synonyms within a cluster in terms of *peripheral concepts* (defined in terms of concepts in the ontology), and stylistic and expressive dimensions. On the syntactic–semantic level, syntactic frames represent how words can be combined with others to form sentences.

Figure 2 depicts the clustered model. It shows how the clusters of the near-synonyms of *error*, *order*, *person*, and *object*, in several languages, could be represented in this model. In the figure, each set of near-synonyms forms a cluster linked to a coarse-grained concept, `Generic-Error`, `Generic-Order`, `Person`, and `Object`, respectively. The rectangles in the top part of the figure are concepts defined in the ontology (and are related by inheritance). Clusters in different languages are shown as groups of interrelated lexical entries that connect to the ontology. Thus, the core denotation of each cluster is the concept to which it points. Within each cluster the near-synonyms are differentiated at the subconceptual/stylistic level of semantics, as indicated by dashed lines that link the words in the cluster. (Note that the actual differences in terms of peripheral concepts, stylistic dimensions, and so on, are not shown in this figure.) Not all words in a cluster need be differentiated, and each cluster in each language can have its own 'vocabulary' of peripheral concepts, stylistic dimensions, and attitudinal dimensions for differentiating its near-synonyms. Note that the figure does not show a representation on the syntactic–semantic level. We can now describe the internal structure of a cluster in more detail starting with two examples.

Figure 3 depicts part of the representation of the cluster of *error* nouns (*error, mistake, blunder, . . .*). It shows the core denotation (the white region), three of its peripheral concepts (the concepts in the shaded regions linked to the core concept by solid arrows), an expressed attitude, and the stylistic dimension of 'concreteness'. (Note that concepts are depicted as regular rectangles, whereas stylistic dimensions and attitudes are depicted as rounded rectangles.) The

[2] It is very probable that synonym clusters could be built automatically by applying statistical techniques, such as cluster analysis, on large text corpora. For instance, Church *et al.* (1994) give some results in this area.

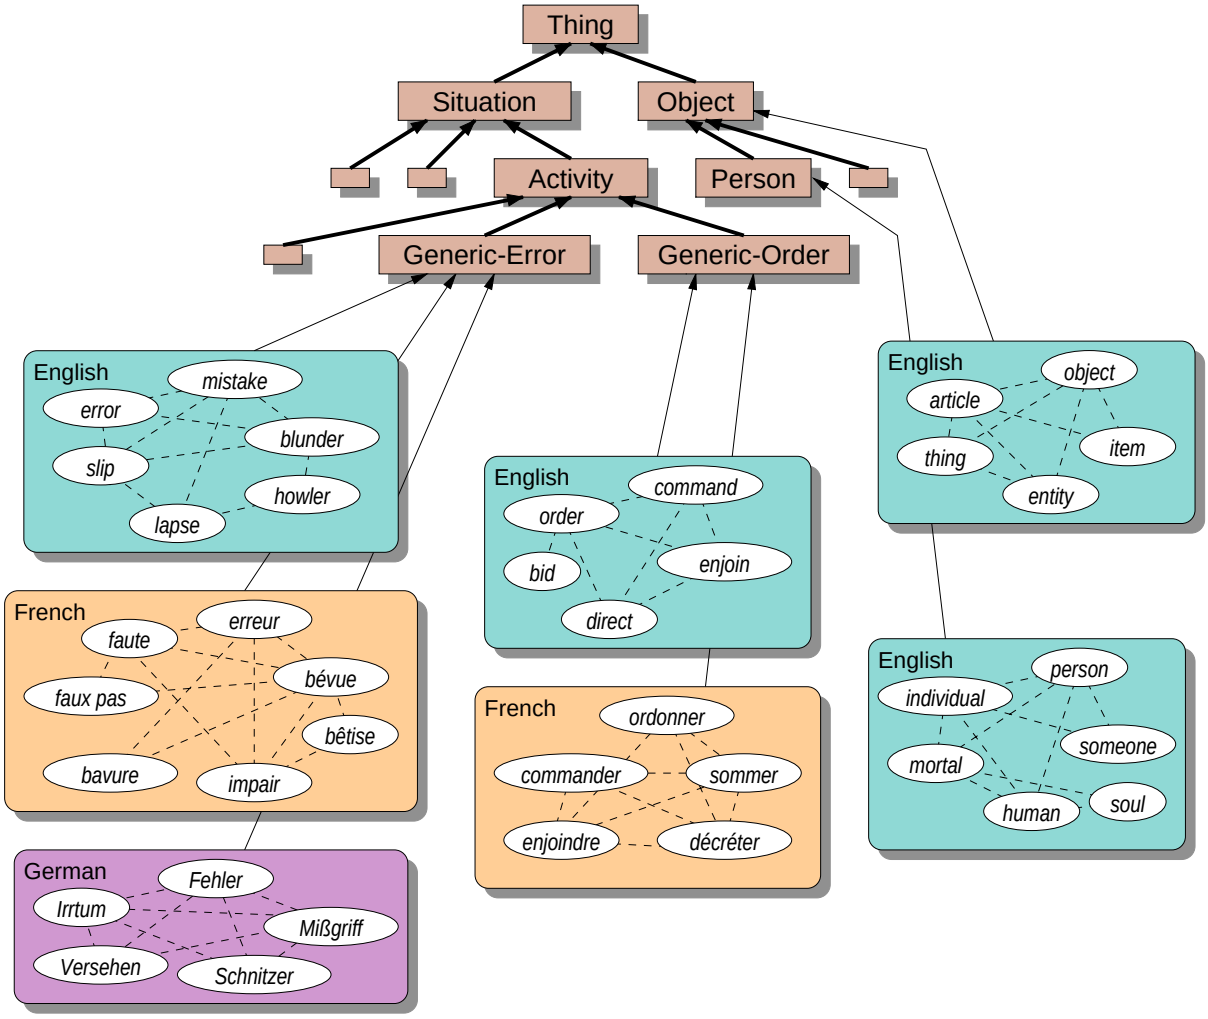

Figure 2: A clustered model of lexical knowledge.

core denotation and peripheral concepts together form a directed graph of concepts linked by relations; the individual concepts and relations are defined in the ontology. In this cluster, the core represents an activity by a person (the actor) that is a deviation. The peripheral concepts represent that a word in the cluster can potentially express, in relation to its near-synonyms, the stupidity of the actor of the error, criticism of the actor (with different levels of severity: `low`, `medium`, or `high`), and misconception as cause of the error. But while all of the near-synonyms in the cluster will convey the concepts in the core denotation, the peripheral concepts that will be conveyed depend on each near-synonym. This is depicted by the two shaded regions (that are bounded by the solid line and the dashed line), which each contain the concepts, styles, and attitudes conveyed by their associated near-synonyms. Thus, *error* can convey a medium degree of severity of criticism compared to the high degree that *blunder* can convey; *error* does not convey stupidity whereas *blunder* can; *blunder* can also express a pejorative attitude towards the actor of the error, but *error* does not express any attitude; and *error* and *blunder* differ stylistically in their degree of concreteness. Notice that the attitude connects to the concept `Person`, because all attitudes must be directed towards an entity of the situation to be expressed. Stylistic dimensions, on the other hand, are completely separate from the graph of concepts. Also, the indirectness of expression of each of the peripheral concepts by each of the near-synonyms is not shown in this diagram, but see below. The appendix gives the complete representation of the cluster in the formalism of our model.

Similarly, figure 4 depicts the cluster of *order* verbs (*order, enjoin, command, ...*), including three of its peripheral concepts and one stylistic dimension. In this cluster, the core represents a communication by a person (the sayer) to another person (the sayee) of an activity that the sayee must perform. The core involves concepts that are not actually

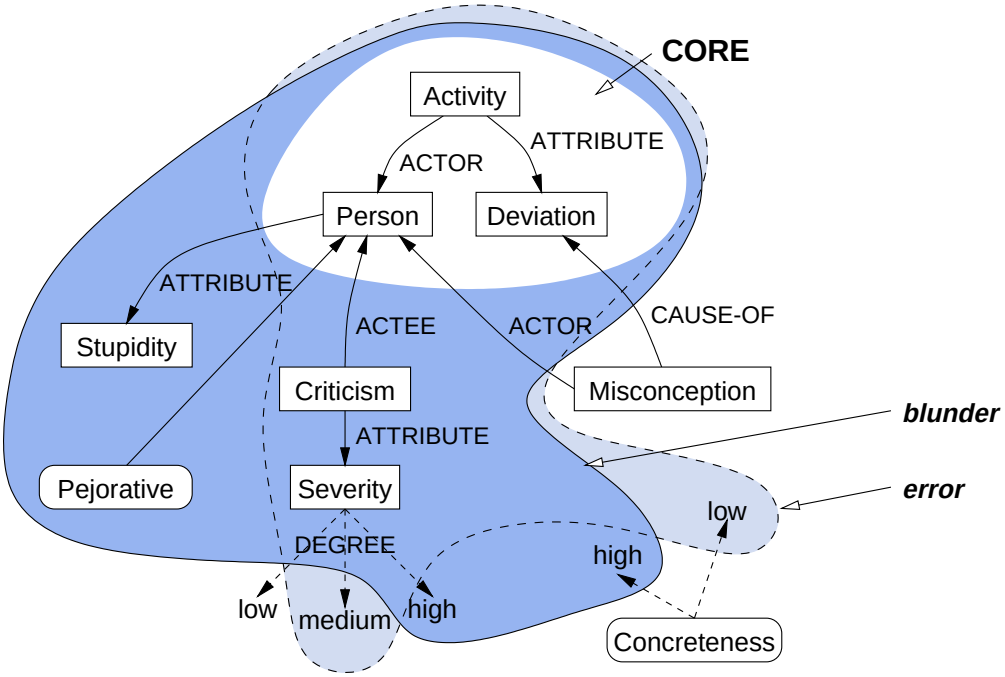

Figure 3: The core denotation and some of the peripheral concepts of the cluster of *error* nouns. The two shaded regions, bounded by the solid line and the dashed line, show the concepts (and attitudes and styles) that can be conveyed by the words *error* and *blunder* in relation to each other.

lexicalized by any of the words in the cluster (e.g., the sayer of the order is not lexicalized by the verb *order*; it must be lexicalized separately). Such concepts, indicated by dashed rectangles, are required in the representation because the peripheral concepts refer to them. The peripheral concepts represent that a near-synonym can express the authority of the sayer (with possible values of `Official` or `Peremptory`), a warning to the sayee, and/or the imperativeness of the activity (with possible values of `low`, `medium`, or `high`). The figure shows the difference between *order* and *enjoin*.

Core denotation

The core denotation of a cluster of near-synonyms is the basic inherent context-independent (and in this formulation of the model, language-neutral) denotation shared by all of the near-synonyms. It links the near-synonyms of the cluster to the ontology, so that it can be used as the necessary applicability condition of all of the near-synonyms in the cluster. (That is, in text generation, it must match part of the input to the system in order for any of the words to be available for choice. In analysis, it is the portion of meaning that should be added to the meaning representation of an utterance.)

A core denotation is represented as a directed graph of concepts linked by relations. The graph can be of arbitrary size, from a single concept (such as `Generic-Error`), up to any number of interrelated concepts (see figures 3 and 4). However, the core denotation must be specified at a sufficient level of granularity that the peripheral concepts can also be specified. So, for instance, in the *error* cluster, it was not possible to use the simple concept `Generic-Error`, because the peripheral concepts of the cluster refer to finer-grained aspects of the concept (the actor and the deviation). We used an equivalent, but finer-grained, representation of the concept. In our formalism, this core denotation is represented as follows:[3]

```
(ROOT Activity
      (ACTOR (V1 Person))
      (ATTRIBUTE (V2 Deviation)))
```

Each node has a variable associated with it so that it can be referred to by other parts of the cluster representation. The `ROOT` variable has a special status: it 'coerces' the graph into a tree-like structure. Without a tree-like structure matching conceptual configurations to one another and to instantiations of configurations (which form the input to a text generation system) would become a potentially complex, even intractable, computational problem.

Peripheral concepts

Informally, peripheral concepts form the basic vocabulary of fine-grained denotational distinctions. For instance, in differentiating the *error* words, a lexicographer would first decide that the basic peripheral concepts required might be

[3] In the formalism, concepts are capitalized and relations are in all capitals, as are variables, but the meaning should be evident from the context.

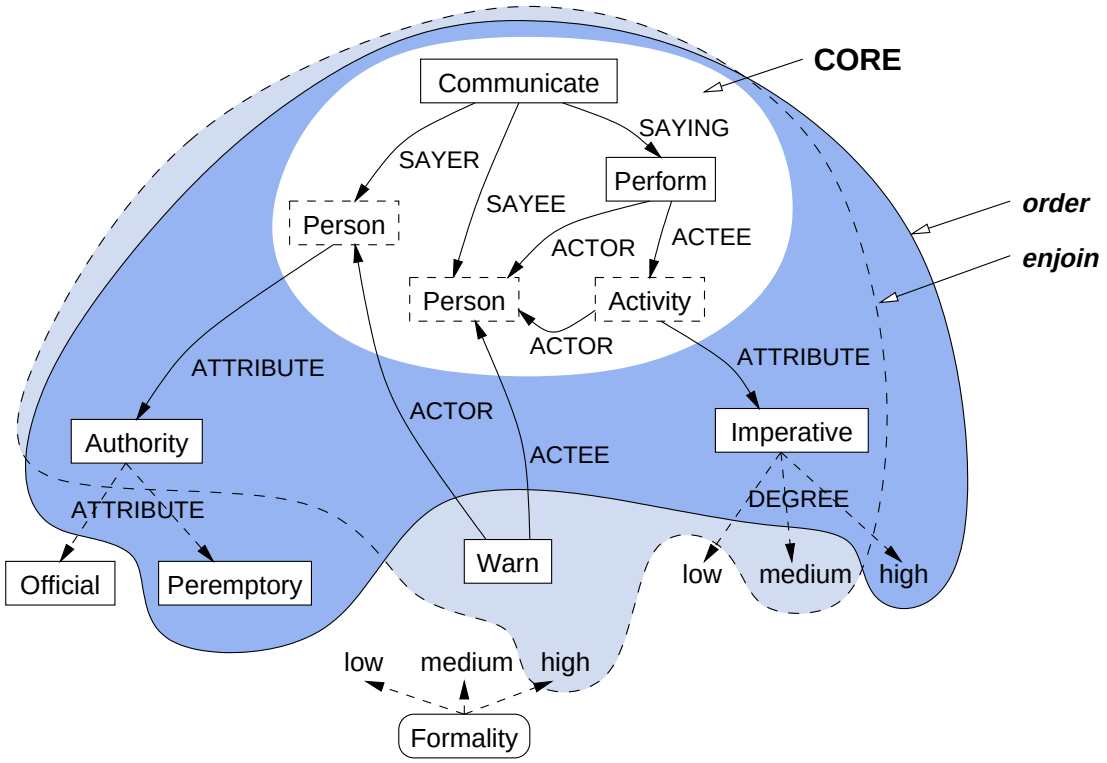

Figure 4: The core denotation and peripheral concepts of the cluster of *order* verbs. The two shaded regions, bounded by the solid line and the dashed line, show the concepts that can be conveyed by the words *order* and *enjoin* in relation to each other.

'stupidity', 'blameworthiness', 'criticism', 'misconception', 'accidentalness' and 'inattention'. Then the lexicographer would proceed to distinguish the near-synonyms in terms of these concepts, for instance, by specifying that *mistake* involves a less severe form of criticism than *error*.

More formally, peripheral concepts are structures of concepts defined in the same ontology as core denotations are defined in. In fact, every peripheral concept in a cluster must 'extend' the core denotation in some way, because, after all, peripheral concepts represent ideas related to the core meaning of a cluster of near-synonyms. Peripheral concepts are represented separately from the core denotation, which gives them special status. They are used to represent non-necessary and indirect aspects of word meaning. That is, they are concepts that might be implied, suggested, emphasized, or otherwise, when a word is used, but not always.

In our formalism, a peripheral concept is, at minimum, a single concept of the ontology with one relation to the core, but as with core denotations they can be of arbitrary size. A few examples for the *error* cluster include the following:

```
(P1 Stupidity (ATTRIBUTE-OF V1))
(P3 Criticism (ACTEE V1)
              (ATTRIBUTE (P3-1 Severity)))
(P4 Misconception (CAUSE-OF V2) (ACTOR V1))
```

Notice that each concept refers to the core denotation using the variables `V1` or `V2`. Each concept has a variable of its own (`P1`, `P3`, or `P4`), to which the denotational distinctions can refer. See the appendix for the complete representation of the peripheral concepts of this cluster.

Distinctions between near-synonyms

We will discuss only denotational distinctions in this paper; representing stylistic and attitudinal distinctions is relatively straightforward in this model.

Although we would like to represent differences explicitly as first-class objects (so that we can reason about them during processing), we can't do so if we are to have a practical and efficient means of representing the differences. Our solution is to represent differences implicitly, but to provide a method for computing explicit differences as needed. Thus, we associate with each near-synonym in a cluster a set of distinctions that are taken to be relative within the cluster; the cluster establishes the local frame of reference for comparing them. So, a word's set of distinctions implicitly differentiates the word relative to its near-synonyms. In other words, if one considers the peripheral concepts to be dimensions, then the set of distinctions situates a word in a multi-dimensional space relative to its near-synonyms.

In our formalism, each denotational distinction refers to a particular peripheral concept. A distinction can encode, in addition to fine-grained differences in the concept itself, differences in the manner in which the concept is expressed

(i.e., differences in indirectness, strength, and frequency). Thus, a distinction is a 5-tuple of the following components: the lexical item (i.e., near-synonym), frequency of expression, strength of expression, indirectness of expression, and concept. Here, we will discuss only differences in the concept, and not in the manner of expression.

Since each peripheral concept is, in effect, a dimension of variation, a distinction actually specifies a value on the dimension. A peripheral concept can be a binary, continuous, or discrete dimension. In figure 3, we represented `Stupidity` and `Misconception` as binary dimensions. Thus, these concepts can either be expressed or not expressed (subject to the manner of expression) by a near-synonym (e.g., *blunder* expresses `Stupidity`; *error* does not). Also in the *error* cluster, we represented `Severity` of `Criticism` as a continuous dimension, which can take either a number (in the range 0–1) as a value, or a range of values specified by a fuzzy set. So, for instance, *blunder* has a value of `high` on this dimension, where `high` represents a range of numeric values at the high end of the dimension. Finally, a discrete dimension, such as `Authority` in the *order* cluster can take other concepts as values, in this case, `Peremptory` or `Official`, although, in general, arbitrary conceptual structures are allowed. The following are examples of each of type of distinction:

```
(blunder_l usually medium implication P1)
(blunder_l always medium implication
              (P3-1 (DEGREE 'high)))
(order_l always medium implication
              (P1 (ATTRIBUTE (Peremptory))))[4]
```

For instance, the second distinction specifies that *blunder* always expresses a medium-level implication of a high degree of severity of criticism. The terms 'always', 'medium', and so on, are defined in the complete account of the model given by Edmonds (1999).

Because we treat peripheral concepts as dimensions, we can compute a degree of similarity (or difference) between two near-synonyms, or between a near-synonym and a representation of a meaning. Edmonds (1999, Chs. 8 and 9) discusses a similarity algorithm that takes into account the similarity of conceptual structures and fuzzy sets. The algorithm can be used in lexical choice in order to determine the near-synonym that most closely matches an input given to a text generation system, or to determine the target language near-synonym that most closely matches a source language word in a machine translation system.

Using the model for lexical choice

Very briefly, we will give one example of how the model can be used in a real text generation system. We developed a system called I-Saurus based on the MOOSE system developed by Stede (1999). An input to I-Saurus consists of an instantiation of domain knowledge that must be lexicalized as a sentence (or other linguistic expression) and a set of preferences to guide the lexicalization process. The purpose of the system is to find and choose the set of words (1) whose core denotations completely cover the set of input concepts, (2) that can be combined into a well-formed plan for a sentence, and (3) that collectively meet as many of the preferences as possible.

For example, the following input:

```
Situation:
  (make1 (ACTOR john1)
         (ACTEE error1))
Preferences:
  (imply (criticize1
           (ACTEE john1)
           (ATTRIBUTE (severe1 (DEGREE 'low)))))
  (disfavour john1)
  (low formality)
```

would output:

John makes a blunder.

because this sentence meets more of the preferences than any other possible sentence (e.g., *John makes a mistake*, *John makes a howler*, *John errs*, ...). That is, the word *blunder* covers the input concept `error1`, strongly meets the preference for disfavouring John, and weakly meets the preference for expressing a criticism of low severity. While *blunder* does not meet the preference for low formality (because the formality for *blunder* is unspecified), no other near-synonym in the cluster meets even two of the three preferences.

The choice process in I-Saurus is efficient, because it can find the appropriate cluster or clusters just as easily as it could find words under a conventional model sans near-synonyms; choosing a near-synonym from a cluster is efficient because there are normally only a small number of them per cluster. The process is also robust, ensuring that the right meaning (at a coarse grain) is lexicalized even if a 'poor' near-synonym is chosen in the end.

Conclusion

By introducing the subconceptual/stylistic level of semantic representation, we have developed a new model of lexical knowledge that keeps the advantages of the conventional model—efficient paraphrasing, lexical choice (at a coarse grain), and mechanisms for reasoning—but overcomes its shortcomings concerning near-synonymy. The subconceptual/stylistic level is more expressive than the top level, yet it allows for tractable and efficient processing because it isolates the expressiveness in small clusters. The model reconciles fine-grained lexical knowledge with coarse-grained ontologies using the notion of granularity of representation. Finally, we are investigating the extension of the model to account for differences across clusters.

Acknowledgments

Our work is financially supported by the Natural Sciences and Engineering Research Council of Canada and the University of Toronto.

[4] `P1` refers to the peripheral concept of `Authority` in the cluster of *order* verbs.

Appendix

The following is the representation of the cluster of *error* nouns in our formalism. Tokens ending in '_l' represent lexical items. Tokens in all capitals are either variables (for cross-reference) or relations, which should be clear from the context. Capitalized tokens are concepts. And tokens in lowercase are values of various features (such as 'indirectness' and 'strength') defined in the model. We did not discuss the 'p-link' and 'covers' fields in the paper, because of lack of space.

```
(defcluster error_C
    ;;; from Gove
    :syns (error_l mistake_l blunder_l slip_l lapse_l howler_l)
    :core (ROOT Generic-Error)
    :p-link ((V1 (:and (Person V1) (ACTOR ROOT V1)))
             (V2 (:and (Deviation V2) (ATTRIBUTE ROOT V2))))
    :covers (ROOT)

    :periph ((P1 Stupidity (ATTRIBUTE-OF V1))
             (P2 Blameworthiness (ATTRIBUTE-OF V1))
             (P3 Criticism (ACTEE V1) (ATTRIBUTE (P3-1 Severity)))
             (P4 Misconception (CAUSE-OF V2) (ACTOR V1))
             (P5 Accident (CAUSE-OF V2) (ACTOR V1))
             (P6 Inattention (CAUSE-OF V2) (ACTOR V1)))

    :distinctions (
     ;;  Blunder commonly implies stupidity.
     (blunder_l usually medium implication P1)

     ;;  Mistake does not always imply blameworthiness, blunder sometimes.
     (mistake_l sometimes medium implication (P2 (DEGREE 'medium)))
     (blunder_l sometimes medium implication (P2 (DEGREE 'high)))

     ;;  Mistake implies less severe criticism than error.
     ;;  Blunder is harsher than mistake or error.
     (mistake_l always medium implication (P3-1 (DEGREE 'low)))
     (error_l always medium implication (P3-1 (DEGREE 'medium)))
     (blunder_l always medium implication (P3-1 (DEGREE 'high)))

     ;;  Mistake implies misconception.
     (mistake_l always medium implication P4)

     ;;  Slip carries a stronger implication of accident than mistake.
     ;;  Lapse implies inattention more than accident.
     (slip_l always medium implication P5)
     (mistake_l always weak implication P5)
     (lapse_l always weak implication P5)
     (lapse_l always medium implication P6)

     ;;  Blunder expresses a pejorative attitude towards the person.
     (blunder_l always medium pejorative V1)

     ;;  Blunder is a concrete word, error and mistake are abstract.
     (blunder_l high concreteness)
     (error_l low concreteness)
     (mistake_l low concreteness)

     ;;  Howler is an informal term
     (howler_l low formality))
)
```

References

Barnett, J.; Mani, I.; and Rich, E. 1994. Reversible machine translation: What to to when the languages don't match up. In Strzalkowsi, T., ed., *Reversible grammar in natural language processing*. Boston: Kluwer Academic Publishers.

Chapman, R. L., ed. 1992. *Roget's International Thesaurus*. HarperCollins Publishers, 5th edition.

Church, K. W.; Gale, W.; Hanks, P.; Hindle, D.; and Moon, R. 1994. Lexical substitutability. In Atkins, B., and Zampolli, A., eds., *Computational Approaches to the Lexicon*. Oxford University Press. 153–177.

Clark, E. V. 1992. Conventionality and contrast: Pragmatic principles with lexical consequences. In Lehrer and Kittay (1992a). 171–188.

Cruse, D. A. 1986. *Lexical Semantics*. Cambridge University Press.

DiMarco, C.; Hirst, G.; and Stede, M. 1993. The semantic and stylistic differentiation of synonyms and near-synonyms. In *AAAI Spring Symposium on Building Lexicons for Machine Translation*, 114–121.

Edmonds, P. 1999. *Semantic Representations of Near-Synonyms for Automatic Lexical Choice*. Ph.D. Dissertation, Department of Computer Science, University of Toronto. *http://www.cs.toronto.edu/compling/Publications/Abstracts/Theses/EdmondsPhD-thabs.html*.

Evens, M., ed. 1988. *Relational Models of the Lexicon: Representing Knowledge in Semantic Networks.* Cambridge University Press.

Frege, G. 1892. Über Sinn und Bedeutung. *Zeitscrift fur Philosophie und Philosophie Kritik* 100:25–50. English translation: On sense and reference. In Geach, P. & Black M., editors, *Translations from the Philosophical Writings of Gottlob Frege*. Blackwell, Oxford, 1960.

Gentner, D., and Markman, A. B. 1994. Structural alignment in comparison: No difference without similarity. *Psychological Science* 5(3):152–158.

Goodman, N. 1952. On likeness of meaning. In Linsky, L., ed., *Semantics and the Philosophy of Language*. Urbana, Ill.: University of Illinois Press.

Gove, P. B., ed. 1973. *Webster's New Dictionary of Synonyms*. G. & C. Merriam Co.

Hayakawa, S. I., ed. 1994. *Choose the Right Word: A Contemporary Guide to Selecting the Precise Word for Every Situation*. New York: HarperCollins Publishers.

Hirst, G. 1995. Near-synonymy and the structure of lexical knowledge. In *AAAI Symposium on Representation and Acquisition of Lexical Knowledge: Polysemy, Ambiguity, and Generativity*, 51–56.

Hovy, E. 1988. *Generating Natural Language Under Pragmatic Constraints*. Lawrence Erlbaum Associates.

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

Katz, J. J., and Fodor, J. A. 1963. The structure of a semantic theory. *Language* 39:170–210.

Kay, M. W., ed. 1988. *Webster's Collegiate Thesaurus.* Springfield, Mass.: Merriam-Webster.

Lehrer, A., and Kittay, E. F., eds. 1992a. *Frames, Fields, and Contrasts: New Essays in Semantic and Lexical Organization*. Lawrence Erlbaum Associates.

Lehrer, A., and Kittay, E. F. 1992b. Introduction. In *Frames, Fields, and Contrasts: New Essays in Semantic and Lexical Organization* (1992a). 1–20.

Lyons, J. 1977. *Semantics*. Cambridge: Cambridge University Press.

Lyons, J. 1995. *Linguistic Semantics: An Introduction*. Cambridge: Cambridge University Press.

Markman, A. B., and Gentner, D. 1993. Splitting the differences: A structural alignment view of similarity. *Journal of Memory and Language* 32:517–535.

Nirenburg, S., and Defrise, C. 1992. Application-oriented computational semantics. In Rosner, M., and Johnson, R., eds., *Computational Linguistics and Formal Semantics*. Cambridge University Press. 223–256.

Pustejovsky, J. 1995. *The Generative Lexicon*. Cambridge, Mass.: MIT Press.

Quine, W. V. 1951. Two dogmas of empiricism. *Philosophical Review* 60:20–43.

Room, A. 1981. *Room's Dictionary of Distinguishables.* Boston: Routledge & Kegan Paul.

Saussure, F. 1916. *Cours de linguistique générale*. Translated by Roy Harris as *Course in general linguistics*, London: G. Duckworth, 1983.

Stede, M. 1993. Lexical choice criteria in language generation. In *Proceedings of the Sixth Conference of the European Chapter of the Association for Computational Linguistics*, 454–459.

Stede, M. 1999. *Lexical Semantics and Knowledge Representation in Multilingual Text Generation.* Boston/Dordrecht/London: Kluwer Academic Publishers.

Tarski, A. 1944. The semantic conception of truth. *Philosophy and Phenomenological Research* 4:341–375.

Ullmann, S. 1962. *Semantics: An Introduction to the Science of Meaning*. Oxford: Blackwell.