Preference Syntagmatics

Patrick Hanks

Masaryk University, Brno

Abstract

This paper compares Yorick Wilks's theory of preference semantics with the evidence of English usage in a large corpus and reports the rationale of a project (in progress) that attaches meanings, not to lexical items, but to contextual patterns, in which each lexical item is normally found. These contexts are based on analysis of a large corpus and stored in a *Pattern Dictionary*. In addition to other influences, this work is partly inspired by Wilks's theory of semantic preferences of the 1970s, but there are significant differences. If meanings are attached to words in context instead of in isolation, the formulas needed to express them can express delicate distinctions without being excessively cumbersome. The Pattern Dictionary provides a resource for reducing lexical ambiguity in texts while maintaining interpretative delicacy. The meaning of a word in an unseen document can be estimated by matching its context to one or other of the normal contexts in the *Pattern Dictionary*, which are themselves explicitly linked to a meaning, called a primary implicature. In the past dozen years corpus analysis has shown with increasing clarity that, although the number of all possible syntagmatic combinations in which each word can participate is vast, and indeed perhaps unlimited, the number of **normal** syntagmatic combinations is manageably small. Examples are given of verb entries from the *Pattern Dictionary*.

Preference Syntagmatics

1. Semantic and Syntagmatic Preferences

In the 1970s Yorick Wilks wrote a series of papers (Wilks 1973, 1975, 1977, 1978, 1980) in which he proposed, in contrast to the theory of selectional restrictions of Katz and Fodor (1963), a system of selectional preferences. He argued that, depending on context, a particular interpretation of a word may be *preferred*, but other interpretations must be *accepted* even if they do not satisfy the preference conditions. Among his examples are the following:

1. The adder drank from the pool.
2. My car drinks gasoline.

The verb ***drink*** prefers an animate agent, which invites the interpretation of ***adder*** in 1 as a snake rather than a calculating machine. However, in 2 no animate interpretation of ***car*** is possible, so any reader (including an AI computation) must *accept* a reading in which an inanimate entity is doing the drinking. Like Sherlock Holmes, the interpreter, having ruled out all alternatives, must accept the only remaining explanation, however implausible, as the correct one.

The distinction between preferred contexts and possible contexts is of the greatest importance for processing meaning in texts[1]. 2 is grammatically well-formed, but it is not very normal or regular: cars don't drink. On the other hand, 1 represents a regularity: adders are animate entities and one of the things that animate entities normally do is drink. The Pattern Dictionary is not interested in whether or not this is a true fact about the world; that point is, more importantly, that it is a syntagmatic statement about English—dependent in part on the frequent recurrence of a particular set of nouns (animate entities) as grammatical subjects of ***drink*** and in part on an ontology that gives a list of animate

[1] Unfortunately, in subsequent research on word meaning Wilks and his colleagues did not build on this distinction. Instead, they focused on the computational manipulation of pre-existing dictionaries. A characteristic of these dictionaries is that they say comparatively little about the syntagmatics of words, while saying much (perhaps too much) about word meaning. Wilks himself now acknowledges (Wilks and Ide 2005) that work using such dictionaries as resources has not solved the so-called "WSD" (Word Sense Disambiguation) problem. However, Wilks and Ide do not go on to draw the conclusion that a different kind of dictionary might be needed, nor do they consider the possibility that the "WSD problem" may be badly formulated, i.e. embody false assumptions, and hence be insoluble. Instead, rather surprisingly, they speculate that lexicographers do not have the right kind of expertise for word sense disambiguation and that any attempt at fine-grained delicacy should be abandoned. In this paper I present the rationale for a new kind of dictionary, a dictionary of syntagmatic preferences, which is explicitly designed to show fine-grained, delicate, empirically well-founded connections between meaning and use in a machine-tractable way.

entities. Syntagmatic regularities such as this can be observed in a large corpus and recorded in a *Pattern Dictionary* (see Hanks and Pustejovsky 2005). If a distinction is made between normal and abnormal usage, then normal uses such as 1 can be interpreted by matching to normal patterns, while the interpretation of abnormal uses such as 2 can be left to second-order semantic computation.

Not only adders but also other animate agents drink. What they drink is a liquid, typically water. In fact, this syntagmatic association is so strong that (even without the additional evidence of ‘pool’) it is easy to supply a highly specific default semantic type for the missing direct object in 2. What did the adder drink? Water, of course. How do we know this? Because in any large collection of general texts a statistically significant association between ***drink*** and ***water*** will show up, especially if the subject of verb ***drink*** is a non-human animate agent.

NLP researchers look to large ontologies such as WordNet to answer questions such as, “What is the set of all animate agents?” and “What is the set of substances that animate agents drink?” In WordNet, the hyponyms under the synset containing the expression ***animate thing*** provide a reasonably exhaustive list.[2] So far so good, but if similar questions are asked about the direct object of the verb ***drink*** and about the attributes of ***gasoline,*** WordNet’s answers are less satisfactory. Attributes of the sense of a word are sometimes enumerated in WordNet as if they were separate senses, while in other cases they are not mentioned at all. ***Gasoline***, for example, is listed in WordNet as a ***fuel*** but not as a ***liquid***. ***Water*** occurs in six WordNet synsets, one of which is glossed as “a clear colorless odorless tasteless liquid” and another as “a fluid necessary for the life of most animals and plants”. The latter sense seems to be the relevant one for our question, but unfortunately it is found only as a hyponym of ***food***. You cannot tell from WordNet that the “fluid necessary for the life of animals and plants" is a liquid. The synset that includes ***water*** as a hyponym of ***liquid*** is not linked to drinkable liquids[3]. To solve problems such as this, a possible strategy might be to lump all the synsets for a given target word together and treat them as a single semantic entity, But it turns out that this tactic is even more unsatisfactory, since it removes necessary constraints: For example, ***water*** is listed in another synset as a synonym of ***urine***. Lumping the synsets for ***water*** together could lead to the unfortunate conclusion that urine is food.

To some extent, these are problems of performance by WordNet, at least some of which could probably be corrected fairly easily. However, there are also underlying problems of principle in WordNet, of which the two main ones are a) that it ignores syntagmatics and b) that it presents attributes (e.g. the drinkableness of water) as separate senses. What is needed is an ontology that links attributes to entities and that shows normal syntagmatics, e.g. that what most ***animate entities*** normally ***drink*** is ***water,*** that ***mammals*** also typically drink ***milk***, and that humans also drink ***manufactured beverages*** of various kinds. It needs

[2] The hyponym list also contains a surprisingly large number of unused terms such as ***eutherian, acrodont,*** and ***pleurodont.*** Presumably, these rare or nonexistent words in WordNet do no harm if no one—neither human nor program—ever looks them up.

[3] The term ***drinking water*** is so linked, but this is unsatisfactory. What the adder drank was not 'drinking water' but 'water'.

to enable a purposive link between ***gasoline*** (British: ***petrol***) and ***motor vehicle***, as well as showing that ***gasoline*** is a ***liquid.*** An ontology such as this is a necessary component of the practical application of preference semantics to processing free text. Such an ontology does not yet exist, but exploratory work is in progress on creating one at Brandeis University (Pustejovsky et al. 2004).

2. WSD Revisited: Reducing Lexical Entropy

Wilks and his colleagues have devoted considerable energies over the years to the problem of word sense disambiguation (WSD). The problem, in a nutshell, is to enable an NLP application to decide, as far as is necessary for successful processing, what an ambiguous word means in a given text. If a word has more than one sense, which sense is activated when the word is used in a particular text, and how can we decide which sense to assign it to, and how can we decide what counts as a sense anyway? The first two of these questions presupposethat each content word in a language has a finite set of distinct, mutually exclusive senses. This supposition seems at first sight to be supported by standard dictionaries which list numbered senses, but in fact it is based either on a misreading of those dictionaries or on a misrepresentation of the language by those dictionaries, or both. The idealizations of competing meanings that are sufficient to trigger appropriate interpretations by human readers disguise the extent to which competing meanings of words overlap. As long ago as 1972 Wilks acknowledged this difficulty very clearly, using the example of ***stake***:

> It is very difficult to assign word occurrences to sense classes in any manner that is both general and determinate. In the sentence "I have a stake in this country", and "My stake in the last race was a pound" is "stake" being used in the same sense or not? If "stake" can be interpreted to mean something as vague as "Stake as any kind of investment in any enterprise", then the answer is yes. So, if a semantic dictionary contained only two senses for "stake": that vague sense together with "Stake as a post", then one would expect to assign the vague sense for both the sentences above. but if, on the other hand, the dictionary distinguished "Stake as an investment" from "Stake as an initial payment in a game or race" then the answer would be expected to be different.[4] So, then, word sense disambiguation is relative to the dictionary of sense choices available and can have no absolute quality about it. – Wilks 1972.

Thanks to corpus evidence, we can now see that the lexical ambiguity of the vast majority of content words is of this kind (overlapping polysemy, often in long chains of 'family resemblances'). Mutually exclusive cases such as ***stake***$_1$ (='post') vs. ***stake***$_2$ (= 'investment'), or ***bank***$_1$ (= 'financial institution') vs. ***bank***$_2$ (= 'side of a river'), are rare. The mutually exclusive cases are only a small part of the problem, but according to Wilks

[4] Splitters in the lexicographic game will point out differences in implicatures: for example, one stands to lose one's stake in a race or game (to the bookmaker or winner), whereas one's stake in one's country cannot be lost in the same way.— PWH

and Ide (2006), these are the only cases that have been successfully disambiguated by computational methods. They go on to say:

> There is rarely a need to make distinctions below the homograph-like level for understanding, human or automated.

This is a counsel of despair. Fine-grained distinctions are indeed fuzzy and they do indeed overlap, but any kind of understanding worthy of the name must face up to the problem, not turn away from it. Understanding, properly so-called, involves being able to compute, on a probabilistic, best-likelihood basis, the meaning or implicatures of woprds in context, on the basis of text input (see below for examples).

Elsewhere, Wilks and Ide comment:

> Some linguists (e.g., Lakoff 1987; Heine 1992; Malt et al. 1999) have proposed that polysemy develops via a chain of novel extensions to previously known senses, each building on its predecessors. This idea, and computational methods for it surveyed and discussed in Wilks and Catizone (2002), follows nicely on from proposals for the generative lexicon proposed by Pustejovsky (1995) and others, but adds the notion that at some point, senses diverge enough to deserve independent representation in the lexicon (either computational or mental).

The “chain-of-novel-extensions” model is plentifully supported by etymological evidence, painstakingly developed in the 19th and 20th centuries by Indo-European philologists (see OED *passim*). However, despite its undoubted philological interest, this model does not tell us how to recognize the point at which senses “diverge enough to deserve independent representation”.

> The problem, of course, is in identifying the point at which two senses become distinct enough to warrant separation for the purposes of NLP (or, for that matter, in dictionaries and the mental lexicon). – Wilks and Ide, *op. cit.*

Lexicographers have long recognized that it is unusual for the senses of a word to “become distinct enough to warrant separation” and that therefore any dictionary entry that implies such a separation is simultaneously both an oversimplification and an overelaboration, written for the convenience of human users, not for NLP. However, corpus lexicographers have also noticed something else, the implications of which are more important than the rarity of mutually exclusive sense distinctions. It is this: despite the undeniable existence of indeterminate examples, the majority of uses of each word, in particular verbs and adjectives, fall into a small number of highly distinctive patterns. The patterns consist of syntagmatic regularities. These syntagmatic regularities consist of argument structures or valencies, with alternations such as active/passive, causative/inchoative, object-drop, etc. Each argument of each predicator is populated in normal usage by a set of words that recur frequently and a more or less open-ended set of words which recur rarely or not at all. The recurring words normally share some common feature, either their semantic type or some other semantic property.

In other words, identifying senses that are mutually exclusive is not the heart of the word-meaning problem. The heart of the problem is identifying the influence of normal context on senses that have only partly diverged, and distinguishing normal contexts from abnormal contexts. With this in mind, the WSD problem can be reformulated. Instead of attempting to match uses of words in free text directly to a dictionary that states word meanings, we can match word uses to the normal contexts of those words. For this to be possible, we need two tools: a dictionary of syntagmatic preferences ("a pattern dictionary") and an empirically well-founded ontology that shows the semantic types and other properties of the words that populate the patterns.

Up to now, attempts to describe semantic values of verb arguments (for example PropBank and VerbNet; see Kingsbury and Palmer 2002) have been based on introspection (notably on so-called “Levin classes”, Levin 1993). These do not stand up well to scrutiny in the light of actual usage. More often than not, they are partial and/or inaccurate. (For examples, see Baker and Ruppenhofer 2002; Hanks and Pustejovsky 2005).

Examples of fine-grained contrasting implicatures (preceded by ‘=’) are the following.

2a. A person fires a gun (= the gun stays where it is)
2b. A person fires a bullet at something (= the bullet moves)

3a. He shook his fist at them (= his own fist)
3b. He shook his hand (= someone else's hand)

4a. A lawyer files a lawsuit (= activates a procedure)
4b. A clerk files some papers (= puts them away)
4c. A reporter files a news story (= submits it for publication)

In principle, NLP ought to be able to deal with such distinctions. Examples such as these have provoked fears of unmanageably large lists. Such fears are groundless. The *Pattern Dictionary* currently contains 15 patterns for the verb ***fire***, 36 patterns for the verb ***shake***, and 14 patterns for the verb ***file***; these are sufficient to represent all normal implicatures of these verbs at the sort of fine-grained level indicated here and are no less manageable than the number of senses in machine-readable dictionaries. Moreover, since the patterns for each verb are (mostly) mutually exclusive, no matter how many patterns there are, no combinatorial explosion arises. Many verbs have only one or two patterns. It is true that light verbs such as ***give*** and ***take*** have over 100 patterns; this is because a) the patterns for phrasal verbs such as ***take off*** are subsumed under the main verb and b) light-verb combinations such as ***take notice, take [[Fact]] into account, take account of [[Fact]],*** and ***take the plunge*** are listed as separate patterns. This is not a problem provided that distinctive criteria for recognizing each pattern uniquely can be given.

Although word senses may overlap (i.e. they share components), patterns generally do not. Patterns are generally mutually exclusive. That is, the relationship between patterns and senses is normally many-to-one: a single sense of a word may be associated with several different patterns. The converse (one pattern mapping to two or more senses) does occur, but rarely. When it occurs (other than as a lexicographer's error), it represents a real ambiguity in the language. (See the discussion of 'he drinks' below).

A checklist of words and their normal contexts is a resource for a variety of applications, including inferencing entailments, text summarization, idiomatic text generation, and machine translation. In such a checklist meanings, implicatures, synonyms, translations, and any other desired features are attached to contexts or ‘patterns’, rather than to words in isolation. Putting a word into a normal context greatly reduces (and often completely eliminates) its lexical-semantic entropy. In isolation, the noun ***orange***, for example, could have any of several meanings. Put it together with the verb ***eat*** and the entropy is reduced if not eliminated. Put it together with the verb ***paint*** [plus a surface or physical object as direct object] and a different meaning of ***orange*** is activated. Put the verb ***eat*** together with ***jealousy,*** as in 5, and a different meaning of ***eat*** is generated.

5. It is Lachlan, eaten up by jealousy, who is plotting against my heir.

The pattern underlying 5 is “[[Emotion = Bad]] eat [[Human]] (up)”, where the completive-intensive particle ***up*** is optional. The semantic role of each argument contrasts with another, more common pattern, “[[Human | Animal]] eat [[Food]]”.

So it is with most words. The number of stereotypical collocations for most if not all verbs is both small enough to be collected in a dictionary and widely used enough to reduce and very often eliminate uncertainty of meaning. We now know that human linguistic behaviour is much more highly patterned than was recognized before the statistical analysis of large and representative collections of text became possible. The pattern dictionary provides patterns of preferred contexts, to which actual contexts can be matched. Of course, there is not always a perfect fit between each actual occurrence and any of the patterns in the pattern dictionary, so a preferential approach must be adopted. Let us return for a moment to 1 and 2 and ask, are both these sentences normal? It is intuitively obvious that animate entities drinking water is normal and cars drinking gasoline is abnormal, and this obvious observation is supported by corpus evidence. It is necessary to distinguish systematically between the two.

The tension between actual context and preferred context in unusual uses of a word might be expected to result in incoherence, but in fact it very often contributes metaphorical resonance to the meaning of the text in which the word occurs. Such metaphorical resonance is evidently part of a word’s linguistic Gestalt—but this must be the subject of a future investigation. The exploitation of semantic and phraseological preferences is clearly rule-governed, but exploitation rules remain to be explicated systematically. The rules governing the exploitation of preferences can only be properly explored once the preferences themselves have been satisfactorily identified. That is the current task of the team building the Pattern Dictionary.

The proposition “Animate entities drink water” at first sight appears to be reminiscent of those found in CYC (Lenat 1990), a large knowledge base of common-sense propositions about terms, their meanings, and their entailments, organized according to theories and microtheories.

6. Fred saw the plane flying over Zurich
7. Fred saw the mountains flying over Zurich

CYC claims that it can distinguish the different meanings of ***flying*** in 6 and 7 because it knows that Fred is a human and humans can fly, while mountains don't. However, CYC is not evidence-based and does not use syntagmatic organization to support its microtheoretical organization of propositions about terms. It is hard to see how the coarse-grained syntactic parser used by CYC could be made more delicate (more fine-grained) without some statistically well-founded form of syntagmatic lexical clustering. Thus the semantic component of CYC is probably overloaded in the case of normal everyday words, though it is hard to be sure. This may be one reason why CYC claims successful applications in the processing of terms only in restricted domains such as business and physics, rather than in general language processing. As far as one can tell, the usefulness of CYC in processing free text is impaired by a combinatorial explosion of the possible senses of common words. If this is true, then syntagmatic filtering by matching free text with normal patterns may be at least part of the solution.

3. Why Standard Dictionaries Won't Do

Despite his 1972 reservations, Wilks and his colleagues devoted considerable efforts to making an English learner's dictionary (*Longman Dictionary of Contemporary English*, LDOCE: Procter et al., 1978) machine-tractable and massaging the information in it for various purposes, including distinguishing the meanings of polysemous words. As Wilks and Ide (2006) make clear, this approach has been at best only partially successful, in part for the reasons stated in the preceding section, but also because essential information about context is not given in standard English dictionaries such as LDOCE. For human users it is not necessary to state the obvious, but for AI applications and linguistic computing it is. So, for example, the LDOCE entry for the verb ***drink*** does not explicitly state that an animate subject is expected. The entries for ***human, animal, beast, animate,*** etc., quite rightly do not explicitly state that humans, animals, and beasts drink. The entry for ***petrol*** says that it is "a liquid used to supply power to the engine of cars and other vehicles", but not that to make a car run someone has to put petrol into the petrol tank, nor that this is a normal meaning of the phrasal verb ***fill up***. The current focus of the *Pattern Dictionary* is on recording salient syntagmatic patterns and their primary implicatures at a sufficiently delicate level to enable understanding of normal text. Having stated the primary implicature of a pattern, it is always possible to add further, secondary implicatures. Having stated all the patterns in which a given word participates, it is also possible to add more, if evidence for them is found—with the proviso that any additional pattern should, ideally, contrast effectively with all other existing patterns for that word.

4. Corpus Analysis and Pattern Identification

Analysis of the phrasal verb ***fill up*** shows that, in normal English, all sorts of containers are filled up with appropriate contents, ranging from ***bottles*** with ***booze*** to ***freezers*** with ***food*** and ***removal vans*** with ***furniture.*** It is therefore necessary to link the verb ***fill up*** to a canonical set of containers. This is done wherever possible by linking the lexical sets in patterns to a semantic ontology. Items that are not normally classified as containers (e.g. ***wallet***) may be coerced to be honorary members of the set of containers when used with ***fill up***, e.g. ***wallets*** may be filled up with ***money*** and ***people*** with ***food,*** but interpretation of Given the mechanism of coercion, it is not necessary to add the lexical items ***wallet*** and ***person*** under the semantic type [[Container]] and indeed it would be wrong to do so. At the same time, ***filling up a car with petrol*** is so conventionalized that it is generally underspecified, i.e. the actual container (the petrol tank) is not mentioned and must be inferred from the context. Part of the job of the Pattern Dictionary is to record explicitly elements of a sentence that may be underspecified in everyday usage (as in 9 below).

It is normally clear from the context (both the physical situation of a speaker and textual collocates such as ***driver***, ***car,*** or ***garage***) that the utterance “Fill her up” means ‘pour petrol into the petrol tank of the car’, not ‘pour petrol into the car (i.e. the passenger compartment)’, still less ‘pour liquid into my female passenger’. In context, speakers regularly underspecify messages. At a petrol station, ***petrol*** does not need to be mentioned, unless it be contrastively (to distinguish it from ***diesel*** fuel), and in English (as opposed to Czech or Russian) people more often than not do not specify that it is the ***petrol tank*** that is filled up. In everyday conversation at petrol stations, the holonymy passes unnoticed and unremarked. Such underspecified utterances trade on established linguistic (not merely situational) conventions. The *Pattern Dictionary* provides a mechanism for supplying the missing semantic values and arguments of such arguments. Some patterns are mutually exclusive, but more often the contrast is between general and particular:

8. [[Human]] fill [[Container]] {up} {with [[Stuff | [PL]Phys Obj]]
9. [[Human]] fill [[Road Vehicle = Petrol Tank]] {up} (with [[Petrol]])

Semantic ontologies record that ***petrol tank*** is a [[Container]] and that ***petrol*** is [[Stuff]]; thus Pattern 9 is merely a subset of Pattern 8. The point of listing them separately is to provide a mechanism for showing semantic correlations for the marked case—here, ***petrol*** and ***car*** or ***tank*** in relation to ***fill up.*** In the pattern dictionary, the specific is always preferred to the general. These two patterns then contrast with further patterns.

Having specified all the normal patterns of usage for a given verb, together with their implicatures, the *Pattern Dictionary* is ready for use. It states primary implicatures (e.g. that Pattern 9 implies that a person inserts the nozzle of a petrol pump into the petrol tank of a car and pumps petrol into the petrol tank) and allows for the addition of any number of secondary implicatures (e.g. that this event takes place at a filling station or that this is necessary in order provide fuel to make the car’s engine operate).

Now let us look at how the *Pattern Dictionary* deals with a more fine-grained distinction.

10. [[Human]] drink [[Liquid = Beverage]]

IMPLICATURE:[[Human]] takes [[Liquid =Beverage]] in through the mouth and swallows it.

11. [[Animate Entity]] drink [[Liquid = Water]]
    IMPLICATURE: [[Animate Entity]] takes [[Liquid = Water]] in through the mouth and swallows it.
    SECONDARY IMPLICATURES:
    [[Animate Entity]] does this in order to quench thirst.
    Doing this is necessary for the survival of [[Animate Entity]]

12. [[Human]] drink [NO OBJ]
    IMPLICATURE:[[Human]] takes alcohol in through the mouth and swallows it.
    SECONDARY IMPLICATURES:
    [[Human]] becomes drunk as a result of doing this.
    [[Human]] regularly becomes drunk.
    Doing this is bad.
    COMMENT: only in simple (not continuous) tenses.

These three patterns illustrate the distinction between lumping and splitting in lexicography. Distinguishing such closely overlapping senses enable delicate semantic implicatures to be represented. A more coarse-grained pattern dictionary (of the kind that seems to be the current goal of PropBank) would settle, perfectly reasonably, for a single pattern, namely 11, without secondary implicatures. (Note, however, that 10 is the most frequent pattern for ***drink*** (such is the anthropocentric nature of human language), while 11 is the least marked and 12 is the most marked).

By distinguishing 12 as a separate pattern from 10 and 11, the following semantic representation can be made. When an antelope or an adder drinks (intransitive), we have a simple object-drop alternation of the transitive verb, whereas when a human drinks (intransitive), there is an ambiguity. It may be a simple object-drop alternation (as when someone at a tea party picks up their cup and drinks), or it may imply getting drink or even habitual intake of alcohol, typically in excessive quantities. Pattern 12 provides a phraseological hook on which to hang the implicature that, if someone drinks, he or she may have alcoholic tendencies. Other textual clues (but not the argument structure, which is the central concern of the Pattern Dictionary) may enable English-language users to distinguish between a single occasion of drinking (in the sense of getting drunk) and habitual drinking (in the sense of being or risking being an alcoholic).

There is a stable set of just over 6000 verbs in normal use in English, depending on whether phrasal verbs are counted separately[5]. An account of all normal lexicosyntagmatic patterns of these verbs is the first target of the *Pattern Dictionary*. Identifying the normal usage patterns of 6000 linguistic items is indeed a large and ambitious undertaking, but considerably less ambitious than compiling a standard dictionary. By contrast, identifying

[5] This figure excludes verbs such as ***abscise, absquatulate, attemper, attorn, auscultate,*** and ***auspicate,*** which are found in WordNet and some dictionaries but not in ordinary usage.

the normal usage patterns of nouns would be a task fifteen or twenty times larger, and would risk missing the interrelatedness of nouns and verbs.

5. Linking Patterns to an Ontology

Lexical items are plugged into a semantic ontology or 'type system'. The model for this the Brandeis Semantic Ontology (BSO) (see Rumshisky et al., 2006). In the BSO meaning elements and the links between them are organized independently of the lexical items that are associated with each type. At the same time, each sense of a verb is associated with one or more patterns, which express the meaning not only of the verb but also of its arguments, in terms of semantic types. Consider one of the patterns mentioned above for the verb ***drink:***

13.[[Human]] drink [[Liquid = Beverage]]

This pattern is linked to three types in the Brandeis Semantic Ontology:

14. [[Animate Entity]] (which has[[Human]] as one of its subtypes)
15. [[Drink Activity]], a subtype of [[Ingest Activity]], a subtype of [[Activity]], a subtype of [[Event]]. This type lists as its typical subject role [[Animate Entity]] and typical object role [[Beverage]]
16. [[Beverage]] = a subtype of [[Liquid]] (a subtype of [[Material Substance]], a subtype of [[Entity]]). This type lists as its telic role [[Drink Activity]].

It will readily be seen that BSO shows that the event [[Drink Activity]] has a preferential association with [[Beverage]] and vice versa. However, it would be wrong to say that[[Human]] has some kind of preferential association with [[Drink Activity]]. Most people do many other things beside drinking. It may well make sense to create a high-level link between [[Animate Entity]] and [[Activity]]—both people and animals do habitually engage in motivated actions (activities), unlike other physical objects such as rocks—but the wisdom of making such additions to the type system remains to be considered.

Under the Type [[Beverage]] a long list of beverages is given, including the subtype [[Alcoholic Beverage]]. BSO does not attempt to list all the words and names—past, present, and future—that denote, have denoted, or could possibly one day denote a human being. Instead, the partial list of humans in semantic roles (*builder, doctor,* etc.) is supplemented by procedures for named entity recognition.

Now consider two coercions (perfectly well-formed and meaningful but abnormal sentences):

17. John drank some gasoline.

Here again, Event=***drink***=[[Drink Activity]]. The Brandeis Semantic Ontology (correctly) does not type ***gasoline*** as a [[Beverage]], but as [[Fuel]]. It should also say[6] that the constitutive of ***gasoline*** is [[Liquid]], the mutual preference mapping of {[[Animate Entity]] <--> [[Drink Activity]]} and {[[Drink Activity]] <--> [[Liquid Substance]]} would provide the correct interpretation (namely that John has engaged in a [[Drink Activity]]) and we could call for the ambulance to take John to the hospital.

18. John's car drinks gasoline.

In 18 the default event type must be coerced to [[Use Activity]]—for the literal meaning of the sentence is that cars use gasoline—so in this case the argument types are not inherited from the Event type, but vice versa. This is done by linking [[Driving Vehicle]] to [[Use Activity]] & [[Fuel]]. The combination [[Driving Vehicle]] plus [[Use Activity]] & [[Fuel]] outweighs the normal interpretation of the verb ***drink*** as a [[Drink Activity]] and coerces it to [[Use Activity]].

6. Refining Preferences by Corpus Analysis

Wilks (1980: p. 141) states that the verb ***fire at*** prefers an animate object, and discusses sentence 19.

19. John fired at a line of stags.

The notion that ***fire at*** has a preference for an ANIMATE object is widely repeated in the literature, but as it happens it is not true. It simultaneously overrestricts and underrestricts. On the one hand, targets that are fired at regularly include VEHICLES and LOCATIONS, not just ANIMATES. On the other hand, the words found in the prepositional object slot (whether ANIMATE or INANIMATE) very often have the semantic property MILITARY (soldiers, tanks, aircraft, defensive positions, etc.). The property MILITARY is not a necessary condition of this argument [7] MILITARY is nevertheless a criterial feature, whereas ANIMATE is not.

If we compare Wilks's example (19) with uses of ***fire at*** in a broad spectrum of general texts such as those of the British National Corpus, we see that *a line of stags* is an outlier in the cluster of terms that populate this slot. Possibly, the situation would be different in a domain-specific corpus of texts, for example, hunting magazines. We shall return to the question of domain-specific semantic values shortly.

Facts such as these present the lexical analyst with a choice, which is resolved by the intended purpose of the lexicon. On the one hand, if the purpose of the lexicon is to express semantic preferences with reasonable certainty, the semantic types should be couched in

[6] At the time of writing the BSO does not say this—probably an oversight.

[7] Examples of prepositional objects of *fire at* with the property HUNTED_ANIMAL, e.g. ***stag***, do of course also occur, and, in a way that is typical of the elastic nature of lexical sets, the set extends outwards to encompass all sorts of PHYSICAL_OBJECTS in addition to the kinds mentioned here.

terms of maximum generality, as in 20. (This is in fact, the position taken by Wilks in the 1970s and subsequently.)

20. [[Human]] fire_at [[PhysObj]]

On the other hand, if the purpose of the lexicon is to enable specific—but necessarily probabilistic rather than certain—predictions to be made about the likely meaning of utterances such as "UNKNOWN fired at UNKNOWN", information about likely semantic roles is needed, as in 21:

21. [[Human (= MILITARY)]] fire_at [[PhysObj (= ANIMATE (= HUMAN) | VEHICLE | LOCATION) & MILITARY]]

A problem for a dictionary based on statistical analysis of corpus data, such as the *Pattern Dictionary*, is that (despite occasional protestations to the contrary) users in the NLP and AI communities, just like any other group of human users, tend to expect that statements in dictionaries represent necessary conditions, not probabilities. Perhaps this expectation can best be countered by adding corpus-based, corpus-monitored statements of observed comparative frequencies automatically to each semantic value, as in 22.

22. [[Human 93% (= MILITARY 25%)]] ***fire_at*** [[PhysObj 94% (= HUMAN 49% | VEHICLE 20% | LOCATION 14%) & MILITARY 26%]]

Informally, this means that 93% of subjects of ***fire_at*** in the sample were identified as Human, and 25% as Military. At the same time, 94% of the direct objects were identified as Physical Objects, subdivided as to 49% Human, 20% Vehicles (*cars, aircraft, tanks, ships,* etc.) , and 14% Locations (military positions, buildings, Israel, etc.). Separately, 26% of the direct objects were identified as having the property Military (*enemy troops, tanks, rebel positions,* etc.). Statistics of this kind can form a serious basis for calculating inferences with measurable probability. Thus, corpus evidence shows that Wilks's "ANI OBJE" for this verb pattern should be replaced (in his terms) by "PHYSOBJ OBJE", since the value ANI here is neither one thing nor the other – neither a strong preference nor a weak probability.

7. Domain-Specific Patterns

Yarowsky et al. (1992) argued that words tend to have a unique interpretation in each document. This claim, known under the catch phrase "One sense per discourse", has been vigorously disputed by Krovetz (1998), who cites data showing Yarowsky's argument (that a very high percentage of words occur in only one sense in any document) is wrong and that, outside Yarowsky's chosen domain of encyclopedia articles, it is quite common for words to appear in the same document bearing different senses. Nevertheless, Yarowsky's point contains a germ of truth. In most documents, at least some of the polysemous content words do have only one meaning, while in other cases the distribution of the different patterns varies according to the domain of the document. We find both domain-specific norms and domain-specific preferences. Thus, the sense of the words in a document is to some extent pre-selected by the domain to which the document belongs, a pre-selection that is reinforced as the document proceeds.

Patterns are invoked in a hierarchy of markedness. To take just one of the groups of norms associated with the verb ***advance*** (the one where it is literally a verb of movement), detailed analysis of a BNC sample shows that the domain is normally military and the grammatical subject (the 'external argument') is an army moving in a particular direction towards an objective[8]. Closely associated with the military norm, and patterned similarly, is a norm in the domain of everyday social behaviour, in which a person ***advances*** on another person. In addition, there are domain-specific norms. In banking and borrowing, ***advancing money*** is the norm; in discourse on philosophical and political topics, ***advancing ideas*** is the norm.

'Normal combinations' must allow for domain-specific preferences; for example, in the subset of BNC composed of reports on soccer matches, the verb ***fire*** exhibits a preference for an adverbial argument with PPs *past the goalkeeper, past the post, into the back of the net*, while in the same subcorpus the verb ***climb*** exhibits a preference for the PP a*bove the defenders*. It is perfectly possible for these soccer-specific preferences to occur in other kinds of texts, but, as a matter of empirically observed fact, they generally don't. Since every text is about something, the role of domain and document in setting lexical preferences needs to be investigated more thoroughly than it has been up to now. Possible but non-preferred combinations become less and less observable as semantic distance from the normal complementation patterns increases. This simple fact underlies all serious work in statistical lexical analysis, from Church and Hanks (1989) to Kilgarriff et al. (2004).

8. The Ubiquity of Preferences

Preferences are everywhere in language systems. Preferences govern both meaning and linguistic behaviour. That is to say, the entire process of constructing an interpretation for a linguistic utterance is governed by a contextual network of interacting preferences. There are no necessary conditions governing the associations between words and meanings, but independent confirmation from different preference systems and indeed from within the same preference system can combine to yield an impression of certainty. The following preference systems at least need to be taken into account:

Lexical preferences: Words prefer the company of certain other words. Example: the verb ***hazard*** overwhelmingly prefers the noun ***guess*** as its direct object, but the role of ***guess*** is also fulfilled, with minor variations in meaning, by other nouns in the same semantic class and by a *that*-clause: *hazard a conjecture, hazard a definition, hazard a suggestion, it seemed sensible to hazard that a man of this standing would have held property in the*

[8] Even though the objective of a military advance is very often left unspecified, it is always implicitly present in the semantics of the verb ***advance*** in this contextual pattern: the difference between an army ***moving*** southwards and an army ***advancing*** southwards is that in the latter case a military objective is implied.

area. It might be tempting to substitute a semantic class, say [[Speech Act]], for the lexical item ***guess*** in the direct-object relation to the verb ***hazard,*** but this would lose specificity of meaning. The verb has a preference for one particular lexical item, namely ***guess***, and that only: more than 50% of uses of ***hazard*** as a verb in both British and American corpora have ***guess*** as their direct object. The verb is indeed used with any of a number of other (non-preferred) nouns denoting a speech act in the direct-object slot, but in the context of this verb, these nouns mean more than just [[Speech Act]]. The strong association between ***hazard*** and ***guess*** colours the interpretation of alternative words such as ***hazard a definition*** and ***hazard a suggestion.*** In this context, even a definition has a guess-like quality. If the direct object slot is broadened to include all kinds of speech acts, the meaning of the phrase becomes blurred or lost.

Syntactic preferences: The verb ***allow*** *prefers* complementation with a direct object and a to-infinitive ("Do not allow a thief to get into your house"), but it *allows* other complementations (e.g. *ditransitive*: "Do not allow a thief access to your house"; *prepositional*: "Do not allow a thief into your house"; *clausal:* "do not allow your house to be accessible to a thief"). Any representation of the underlying syntactic and semantic regularity here must represent both the fact that ***allow*** governs an Event as its argument and the fact that the Event is normally expressed as a to-infinitive.

Domain-specific preferences: As indicated above, for many words the preferred complementation varies from domain to domain: the verbs ***fire*** and ***climb*** have a preferred sense and pattern in soccer journalism that is quite different from the normal sense in other domains. In financial journalism ***climb*** more often refers to the rising value of stocks and shares than to going up a mountain. These facts, too, need to be taken into account in any representation of the lexicon.

9. Conclusions

The main argument of this paper is that the so-called "WSD problem" has been formulated in such a way as to guarantee its own insolubility. The assumption that each word has a finite list of meanings that can be "disambiguated" is plausible but wrong. Meanings cannot be assigned effectively to words in isolation, and that is the source of the WSD problem. The problem needs to be re-formulated before it can be resolved. The reformulation proposed here relies on the fact that each word is associated with an observable, measurable set of normal contexts. In this way, a distinction between normal and abnormal contexts can be established.

This brings us back to Wilks's theory of preference semantics, but instead of trying to formulate word meanings directly, the alternative proposed here is to measure the syntagmatics first—i.e. to compute similarity of the textual context surrounding a word in a document to the best match in an inventory of normal contexts in the *Pattern Dictionary*—and then to select the meaning given for the best-match syntagmatic pattern. The meanings may be more or less delicate, according to a) practicalities, e.g. the needs of the intended application; and b) the constraints imposed by the language itself on what can be recorded in the *Pattern Dictionary*.

Argument structures are expressed in the *Pattern Dictionary* as far as possible in terms of the semantic types and attributes of the Brandeis Semantic Ontology, which states the semantic type of each sense of the word in itself, together with certain of its properties (for example, its 'telic' or purpose).

Given adequate pre-processing (part-of-speech tagging, parsing, and resolution of pronominal anaphora), the pattern of the arguments is very often sufficient to compute the fine-grained meaning of a word in a document with reasonable accuracy. Occasionally, however, it is necessary to invoke additional evidence, such as the domain or text type of the document or, very occasionally, the wider context.

References

Baker, Collin F., and Josef Ruppenhofer. 2002. 'FrameNet's Frames vs. Levin's Verb Classes' in J. Larson and M. Paster (eds.), *Proceedings of the 28th Annual Meeting of the Berkeley Linguistics Society*, pp. 27-38.

Church, K., and P. Hanks. 1990. 'Word association norms, mutual information, and lexicography' in *Computational Linguistics* 16 (1).

Hanks, P., and J. Pustejovsky. 2005. 'A Pattern Dictionary for Natural Language Processing' in *Revue française de linguistique appliquée* 10 (2).

Katz, J. J., and J. A. Fodor. 1963. 'The Structure of a Semantic Theory' in *Language* 39.

Kilgarriff, A., D. Tugwell, P. Rychly, and P. Smrz. 2004. 'The Sketch Engine' in *Proceedings of Euralex, Lorient, France, July 2004,* pp. 105-116.

Kingsbury, P., and M. Palmer. 2002. 'From Treebank to Propbank', *LREC-02: Third International Conference on Language Resources and Evaluation*. Las Palmas.

Krovetz, R. 1998. 'More than One Sense per Discourse'. NEC Princeton NJ Labs. Research Memorandum.

Lenat, D., and R. V. Guha. 1990. *Building Large Knowledge-Based Systems: representation and inference in the Cyc project*. New York: Addison-Wesley.

Levin, B. 1993. *English Verb Classes and Alternations.* University of Chicago Press.

Procter, P., et al. 1978. *Longman Dictionary of Contemporary English*, 1st edn. Harlow: Longman.

Pustejovsky, J., and P. Hanks. 2001. 'Very Large Lexical Lexical Databases'. *ACL Workshop, Toulouse.*

Pustejovsky, J., P. Hanks, and A. Rumshisky. 2004. 'Automated Induction of Sense in Context'. *COLING 2004.* Geneva, Switzerland.

Rumshisky, A., C. Havasi, and J. Pustejovsky. 2006. 'Constructing a Corpus-based Ontology using Model Bias'. *19th International FLAIRS Conference*.

Wilks, Y. 1972. *Grammar, Meaning, and the Machine Analysis of Language.* London: Routledge.

Wilks, Y. 1975a. 'Preference Semantics' in E. L. Keenan (ed.), *Formal Semantics of Natural Language*. Cambridge: Cambridge University Press.

Wilks, Y. 1975b. 'A Preferential, Pattern Seeking Semantics for Natural Language Inference' in *Artificial Intelligence* 6.

Wilks, Y. 1977. 'Good and Bad Arguments about Semantic Primitives' in *Communication and Cognition*, Vol. 10 (3/4).

Wilks, Y. 1978. 'Making Preferences More Active' in *Artificial Intelligence* 11 (3); reprinted in N. V. Findler (ed., 1979), *Associative Network*s. New York: Academic Press.

Wilks, Y. 1980. 'Frames, Semantics, and Novelty' in D. Metzing (ed.), *Frame Conceptions and Text Understanding.* Berlin, New York: De Gruyter.

Wilks, Y., B. Slator, and L. Guthrie. 1996. *Electric Words: Dictionaries, Computers, and Meanings.* Cambridge, MA: MIT Press.

Wilks, Y., and N. Ide. 2006. 'Making Sense about Sense'. Chapter 3 of E. Agirre and P. Edmonds (eds.), *Word Sense Disambiguation: Algorithms and Applications*, Berlin: Springer Verlag.

Yarowsky, D., W. Gale, and K. Church. 1992. ' One Sense Per Discourse' in *Proceedings of the 4th DARPA Speech and Natural Language Workshop*.

Web Sites for Lexical Resources

Corpus Pattern Analysis: http://nlp.fi.muni.cz/projects/cpa/

CYC: http://www.cyc.com

FrameNet: http://framenet.icsi.berkeley.edu/

Propbank: http://www.cis.upenn.edu/~mpalmer/project_pages/ACE.htm

VerbNet: http://www.cis.upenn.edu/~kipper/VerbNet/

WordNet: http://wordnet.princeton.edu/