

Corpus-based methods and cognitive semantics: The many senses of *to run**

Stefan Th. Gries

Abstract

The first major part of this paper is a comprehensive cognitively-oriented analysis of the senses and their interrelations of the verb *to run* along the lines of much recent cognitive work on polysemy. In the second major part, all occurrences of *to run* from the ICE-GB and the Brown Corpus are coded for a variety of linguistic parameters (so-called ID tags), yielding a complete behavioral profile of this verb. On that basis, the paper then discusses several case studies of how such corpus-linguistic quantitative methods can provide objective empirical evidence suggesting answers to some notoriously difficult problems in cognitive linguistics; these include the issue of prototype identification, the (degree of) sense distinctness, the structure of the hypothesized network as well as possibilities of automatic sense identification.

Keywords: polysemy; word sense (disambiguation); behavioral profile; semantic network; cluster analysis.

1. Introduction

The present paper is concerned with word senses from the perspective of cognitive linguistics on the one hand and corpus-linguistics as well as corpus-based lexicography on the other hand. While many recent cognitive-linguistic approaches to polysemy have concerned themselves with polysemous words as network-like categories with many interrelated senses (with varying degrees of commitment to mental representations), corpus-linguistic approaches have remained rather agnostic as to how different word senses are related and have rather focused on distributional characteristics of different word senses. This paper attempts to bridge the gap between these two approaches by demonstrating how cognitive linguistics can benefit from methodologies from corpus linguistics and computational

linguistics; it is therefore a plea for more corpus linguistics in cognitive linguistics and structured as follows: Section 2 provides a by necessity very brief overview of cognitive-linguistic approaches towards polysemy and some of their weaknesses (cf. Section 2.1) as well as some corpus-based approaches (cf. Section 2.2). The review can of course not do justice to the large number of studies on polysemy and especially word sense disambiguation; it merely serves to discuss how the problems of identifying the different senses of a polysemous word have been addressed. Section 3 discusses the senses of the highly polysemous English verb *to run* on the basis of British and American corpus data. Section 4 constitutes the central part of this study. It introduces and exemplifies a few methodologies which increase the descriptive adequacy of cognitively-oriented analyses of lexical items as well as resolve some notoriously difficult questions within the cognitive paradigm. Finally, Section 5 concludes with some further extensions.

2. Distinctions between senses and the relations between them: A short review

2.1. Cognitive-linguistic approaches

One of the central areas of research within cognitive linguistics has been the investigation of polysemy of lexemes and constructions. Traditionally, the idea that a word is polysemous entails that the particular lexeme under investigation (i) has more than one distinct sense (otherwise the lexeme would be considered vague) and (ii) that the senses are related (otherwise the lexeme would be considered homonymous).¹

The former point is usually made on the basis of a variety of well-known ambiguity tests including the logical test, the linguistic (*do so*) test and the definitional test (cf. Geeraerts [1993], Cruse [1986] and Kilgarriff [1997] for detailed discussion). However, these tests often yield mutually contradictory results, which is why cognitive linguists have often posited a continuum of semantic distinctness ranging from clear cases of homonymy on the one hand to clear cases of vagueness on the other hand; cases of polysemy were then located somewhere between these two extremes (cf., e.g., Tuggy [1993] or Croft [1998]). Thus, the distinctness of different senses of a lexeme is considered a matter of degree. Although it is probably fair to say that cognitive linguists have focused on the analysis of how dif-

ferent senses of a word are related to each other, they have of course also been aware that the motivation of sense distinction is a non-trivial issue since the links between senses can only be discussed once the distinctness of senses has been established. Thus, a variety of different approaches have been proposed to deal with this problem; let us briefly consider some examples.

Consider, as a first example, some early studies such as Brugman (1981), Norvig and Lakoff (1987), Lakoff (1987) and Brugman and Lakoff (1988). On the basis of intuition data, nearly every usage event minimally different from another one constitutes a different sense. For instance, Brugman and Lakoff argue that “a polysemous lexical item is a radial category of senses” (1988: 478) and they posit different schemas of the English preposition *over*, which often differ only with respect to properties of the landmark. For instance, in (1a) the landmark (the hill) is vertical whereas, in (1b), it (the yard) is not (Brugman and Lakoff’s [1988: 482–483] examples).

- (1) a. *The plane flew over the hill* → schema 1 (above and across):
vertical extended landmark, no contact
b. *The bird flew over the yard* → schema 1 (above and across):
non-vertical extended landmark, no contact

This so-called full-specification approach (cf. Lakoff 1987) has been criticized for its methodological vagueness (resulting in the high degree of granularity – i.e., minimally different senses – pointed out above), its vagueness of representational convention and its lack of clarity concerning the linguistic and cognitive status of its network architecture (cf. Sandra and Rice [1995] for discussion and exemplification), and other approaches have been adopted to resolve this question on a principled, non-arbitrary basis. For example, Sandra and Rice (1995) as well as Rice (1996) argue in favor of (prepositional) polysemy on the basis of different experimental results. As another alternative, Tyler and Evans (2001) develop a principled-polysemy approach in which a distinct sense of *over* is only posited iff the meaning of *over* in one utterance involves a different spatial configuration from *over* in another utterance and cannot be inferred from encyclopedic knowledge and contextual information.²

However, not all these approaches are equally useful. For example, it is unclear whether the results of the sorting tasks of Sandra and Rice (1995) or Rice (1996) can actually be attributed solely to semantic differences of

the uses (which also undermines the results' utility in refuting monosemy approaches): unlike recent experimental work by, say, Klein and Murphy (2001, 2002), the experimental sentences were not balanced with respect to all lexical items contributing to subjects' decisions. Moreover, different distance measures and clustering algorithms result in different amalgamation schedules and different degrees of granularity, but Sandra and Rice do not provide such details, which makes the evaluation of their findings difficult.

It is only very recently that cognitive linguists have turned to corpus data as a source of evidence for sense distinctions. For example, Croft (1998: 169) argues in favor of investigating the distinctness and conventionality of senses corpus-linguistically. He points out how semantically different direct objects of *to eat* correlate with uses distinct in terms of the arguments they occur with. In addition, Fillmore and Atkins's (2000) discussion of *to crawl* is cognitive-linguistic in the sense that the relations between different senses of *to crawl* are motivated both experientially and frame-semantically, but also truly corpus-based as it relies on an exhaustive analysis of a complete concordance. Finally, Kishner and Gibbs (1996) (as well as Gibbs and Matlock [2001]) discuss associations (of unmentioned strengths) of different senses of the English adverb *just* and *to make* on the one hand to different R1 collocates (i.e., words at the first slot to the right of the word of interest) and syntactic patterns on the other hand. They demonstrate "that people's choice of a sense of *just* is in part determined by the frequency of co-occurrence of particular senses of *just* with particular classes of words" (1996: 27–28) as well as situational characteristics, which results in some resemblance to a frame-semantic approach. Lastly, they propose that such results generalize to (words of) other syntactic categories, e.g. the verb *to run* and, in Gibbs and Matlock (2001: 234), argue that "if polysemous words are best described in terms of lexical networks, then our findings suggest the need to incorporate information about image schemas and lexico-grammatical constructions in drawing links between different senses of a polysemous word", a proposal to which we will return.

2.2. Corpus-based approaches

Especially the last approach by Kishner and Gibbs bridges the gap between cognitively oriented approaches and the linguistic paradigm in which the question of how to determine whether two uses of a particular word instan-

tiate two different senses or not has probably received most attention, namely (corpus-based) lexicography; we will turn to this approach now.

Organizing and formulating a dictionary entry for a word requires many decisions as to whether two citations of a word instantiate senses differing enough that the word's entry needs to be split or whether the citations instantiate senses similar enough to be lumped together. Although the lexicographer's interest in sense distinctions need not coincide with that of linguists of a more theoretical persuasion, the basic question of course remains the same. Given these questions, recent lexicographic work has arrived at the conclusion that word senses as conceived of traditionally do not exist and has therefore adopted an increasingly corpus-based approach. For example, Kilgarriff (1997: 92) argues in favor of "an alternative conception of the word sense, in which it corresponds to a cluster of citations for a word". In the simplest possible conception, "corpus citations fall into one or more distinct clusters and each of these clusters, if large enough and distinct enough from other clusters, forms a distinct word sense" (Kilgarriff 1997: 108). According to him, much lexicographic work more or less conforms to the following characterization: first, call up a concordance for the word. Then, divide the concordance lines into clusters which maximize intra-cluster similarity and minimize inter-cluster similarity. Third, for each cluster, identify what makes the member of a cluster belong together (and change clusters where necessary), and finally, encode these conclusions in lexicographese (cf. also Biber [1993] and Hanks [1996: 82]). Similarly, Hanks (2000: 208–210) argues for a focus on separate semantic components (jointly constituting a word's meaning potential), which can be weighted in terms of their frequency and predictive power for regular word uses.

However, the above is only a very abstract idealization of the actual cognitive processes underlying sense identification and distinction. This and the fact that many of these processes result in apparently subjective decisions is immediately obvious once a user consults different dictionaries on the same word (cf. Fillmore and Atkins [2000] or Gries [2001, 2003a] for discussion). Therefore, corpus-based lexicographers have begun to formulate strategies to provide a more objective foundation for resolving such issues by, for instance, identifying corpus-based traces of meaning components etc. In order to bring together both cognitive-linguistic and corpus-based lexicographic approaches, it is necessary to briefly review the two lexicographic approaches upon which the present approach relies most.

First, Atkins (1987) discusses what she refers to as “ID tags”, i.e. “syntactic or lexical markers in the citations which point to a particular dictionary sense of the word” (Atkins 1987: 24). ID tags are distinguished depending on (i) whether the presence of a particular clue is categorically or probabilistically associated with a particular sense and (ii) whether they testify to a characteristic of the word under investigation directly or indirectly (i.e. via the properties of other words). Atkins then investigates 441 citations of the word *danger* with respect to these ID tags: the word class of *danger* in the citation and more fine-grained distinctions within the word class (e.g. number, countability etc.); the complementation pattern associated with *danger* in the citation (e.g. [NP the danger [PP to [NP health]]], etc.); the function of the phrase in which *danger* occurs (e.g. subject, direct object, complement etc.); and the collocates in a window of ± 7 words.³

Even without a full statistical analysis, Atkins obtains several useful ID tags. For example, the senses of *danger* that can be paraphrased as ‘unsafeness/riskiness’ and ‘someone/something posing a threat’ are associated with, among others, the ID tags in (2) and (3) respectively.

- (2) uncountable noun with no support (as in *They are, however, fraught with danger*)
- (3) countable noun followed by a PP with *in* without a *that*-clause (as in *There are, he agrees, real dangers in a partisan Civil Service*)

It turns out that the predictive power of some ID tags is fairly high, indicating that the ([semi-]automatic) allocation of citations to senses can be further improved; the approach is thus a forerunner of similar work on the automatic identification of semantic roles by Gildea and Jurafsky (2001).

The second lexicographic approach relevant to the present approach is that of Hanks (1996). He argues that the semantics of a verb are determined by the totality of its complementation patterns (1996: 75, 77) and proposes to analyze the usage of a particular word on the basis of the word’s behavioral profile, which basically corresponds to a set of Atkins’s (1987) ID tags together with semantic role generalizations. The different senses of words can then be derived from (i) different patterns within a behavioral profile and (ii) the process of triangulation, i.e. the identification of correlations between two or more lexical sets in different slots associated with the verb. Like Atkins, Hanks does not provide data on the predictive power of the behavioral profile of a verb or the different lexical sets, but he does state that many verbs exhibit strong frequency asymmetries of particular patterns and senses that can aid sense identification considerably.

As we have seen, a few cognitive-linguistic studies (most notably Kishner and Gibbs [1996]) have devoted their attention to how word senses correlate with a narrow range of formal characteristics such as complementation patterns, i.e. what corpus linguists have referred to as colligations. Section 3 first provides a cognitively-oriented polysemy analysis of the English verb *to run* to first of all determine its inventory of senses. The set of applications presented in Section 4 extends Kishner and Gibbs's (1996) hypothesis (that colligations similar to the ones discussed for *just* can be found for other word classes) by taking seriously the notions of behavioral profile and triangulation using the set of lexico-grammatical ID tags employed by Atkins (1987). That is, Section 4 outlines several case studies relating *to run*'s cognitively motivated senses to its corpus-based behavioral profile.⁴

3. *To run*: A cognitively-oriented analysis

In this section, I will discuss, and provide the token frequencies of, the different senses of all 815 instances of *to run* from the British component of the International Corpus of English (ICE-GB; n=391) and the Brown Corpus of American English (n=424).⁵ The senses were identified manually and mainly on the basis of the match of the citation to senses listed in dictionaries and in WordNet 1.7.1.^{6,7} While I will also be concerned with how the different senses of *to run* are related, I will follow Kishner and Gibbs (1996) as well as Fillmore and Atkins (2000) and refrain from elaborating in detail on all cognitive mechanisms relating the different senses, restricting myself to a less rigorous characterization; all examples are taken from the corpus data.

3.1. Intransitive uses of *to run*

The central, or prototypical, sense of *to run* appears to be that of 'fast pedestrian motion' as in (4); cf. Section 4.1 below for a justification of why this sense is considered prototypical.

(4) *Simons had run down to the villa to get help* [n_{of this sense in corpus}=203]

Other, closely related senses are exemplified by (5) (where motion is still fast but not necessarily pedestrian) and (6) (where the motion even need not be fast anymore) – in this example, however, the sentence also implies that the boat makes this journey regularly, a semantic feature we will find again.

(5) *Yet they keep running from one physician to another* [n=4]

(6) *There are three boats that run from the mainland to the Island* [n=24]

Two senses that are closely related to the sense(s) exemplified in (5) and (6) are ‘to move away from something dangerous/unpleasant’ and ‘to move away to engage in a romantic relationship’ in (7a) and (7b); actually, the two senses are similar enough to be considered a single sense provisionally labeled ‘to escape’. Similarly closely related to the central sense is the sense of ‘to look after’ in (7c). These three senses of *to run* need not, but typically do, invoke literal fast pedestrian motion.

- (7) a. *When he loses his temper with her she runs off, taking young Jacob with her* [n=28]
 b. *If Adelia had felt about someone as Henrietta felt about Charles, would she have run away with him?* [n=4]
 c. *At an age when they might want to take things easy <, > many women like sixty-three-year-old Eileen Allen are running around after older relatives* [n=1]

A similarly close relationship to the senses in (5) to (7) is exhibited by the senses ‘motion without control/restraint’ and ‘to meet (unexpectedly)’ in (8) and (9) respectively.⁸

(8) *Dogs ran about, getting in people’s way*

(9) *On my way to the elevator, I ran into Pete*

The sense in (9) can also be extended metaphorically to yield the sense of ‘to speak continuously’ in (10a) with a metonymic understanding of *the bench* and non-human subjects as in (10b); cf. also (23) and (42) below.

- (10) a. *the bench, which numerous times rebuked the Attorney General for letting his witnesses run on* [n=6]
 b. *Then a wild thought ran circles through his clouded brain*

(10) is more remote from the central sense(s) than (8) and (9) – not involving pedestrian motion – but other cases are even more remote: (11a) predicates motion of liquids (i.e. ‘to flow’) and (11b) refers to the potential result of liquid mixing with color (i.e. ‘diffusion of color’).

- (11) a. *The tears ran down my face* [n=16]
 b. *Colors on the towels had run* [n=2]

Then, by some version of the *swarm* alternation (Levin 1993: 53–55) or by what Norvig and Lakoff (1987: 198) have termed profile shift, (12) denotes ‘to exist in abundance’ (cf. Fillmore and Atkins [2000: 103] on the similar *The kitchen was crawling with cockroaches*).

- (12) *Baker, you will have the streets of our American cities running with blood on registration day* [n=1]

Still on the basis of (11) and a profile shift from the liquid to the container out of which it moves, (13) denotes what happens when there is liquid in abundance in a container, namely ‘to overflow’.⁹

- (13) *So when the Big House filled up and ran over, the sisters-in-law found beds for everyone in their own homes* [n=1]

Finally, (14a) refers to the result of a liquid moving out of a container, namely ‘to become used up’, which is extended metaphorically to the domain of time in (14b).

- (14) a. *It has a shelf life of 100 years and will write for three miles before the ink runs out* [n=14]
 b. *Time is running out*

The senses in (14) can also undergo an alternation where the thing becoming used up is demoted from subject status to that of a prepositional complement headed by *out of* and where the former possessor’s role is profiled and becomes the subject of *to run out of*, cf. (15a). If the location from which the liquid is moving is conceptualized as causing the liquid’s motion, we find the sense ‘to emit liquid’ as in (15b).

- (15) a. *we're running out of tea bags* [n=18]
 b. *She can't tell whether her left nostril is running* [n=1]

Another metaphorical extension of *to run* is used as meaning 'to be in charge of something'. Usually, this sense involves the transitive construction (cf. Section 3.2), but the intransitive use in (16) comes about by a metaphor CHANGE OF STATE IS CHANGE OF LOCATION and the metonymy relating organizations to the people who are part of the organization.

- (16) *Thus, if corporations are not to run away with us, they must become quasi-governmental institutions* [n=1]

If the subject is not human but, as in (17), refers to an idea/some information, *to run* can be used to mean that the human referent of the prepositional complement 'fails to control' or 'fails to understand' the subject's referent (via the metaphors IDEAS ARE ENTITIES and UNDERSTANDING/LEARNING IS GAINING PHYSICAL CONTROL OVER AN ENTITY).

- (17) *This [finding out the person some description is referring to and what her surname is] is running away with me* [n=23]

In many other instances, the relation to the central sense is similarly less direct. For example, there are different kinds of what Langacker (1987: 168–173) has called abstract motion, i.e. there are instances where the subject of *to run* is still human, but what is denoted is not literal (pedestrian) motion, but 'metaphorical motion' as in (18a) and (18b) (cf. STATES ARE LOCATIONS), there is the sense of 'to be/to become' in (18c), and the sense of 'to deteriorate' in (18d) (cf. GOOD IS UP).

- (18) a. *we may conceivably run into trouble here* [n=8]
 b. *He ran into the rapture of the depths* [n=1]
 c. *Chief Bob Moore looked his same hick-self; a man mountain running to lard in his middle-age* [n=4]
 d. *pansy seeds, he told me, soon "run down"* [n=3]

Uses similar to the one exemplified in (18d) are also found in more specific contexts, where *to run down* refers not only to 'to deteriorate', but to the slightly different sense of 'to lose power/efficiency', cf. (19).

(19) *Foster listened with [...] patience until Digby ran down, [...]* [n=1]

Similar cases of abstract motion are ‘to check/to rehearse’ in (20a) and (20b) and ‘to campaign’ in (21) (via GOALS ARE DESTINATIONS).¹⁰

(20) a. *Russ ran through the bills* [n=14]
 b. *Cosmo ran through a few historical dates in his mind*

(21) *[...] when Bush was running for the White House, [...]* [n=28]

Then, by analogy to (10) and (17), in instances such as (22) there is not even a human agent performing the (metaphorical) motion.

(22) *Their policy ran counter to the traditional idea that a good fighter was usually a libertine* [n=21]

To run can also be used like a copular verb to denote ‘to have a particular wording’ as in (23) (cf. Langacker 1987: 168–169).

(23) *“[...] Say he is a horse thief”, runs an old adage* [n=12]

(24) exemplifies an image-schema transformation (and Levin’s [1993: Section 4.7.7] *meander* verbs), namely the sense ‘to extend spatially’.

(24) *Street car tracks run down the center of Pennsylvania* [n=55]

This sense has been extended (via TIME IS SPACE) to ‘to continue (to exist) for a certain time period’ (cf. [25a] and [25b]) and to the quantitative senses ‘to amount to’ and ‘to surpass’ in (26) and (27) respectively.

(25) a. *It [a play] ran until past one o’clock* [n=22]
 b. *the diplomatic process has run its course*

(26) *[...] the number may run into tens of millions* [n=14]

(27) *Sales of TV sets at retail ran ahead of the like months of 1959* [n=3]

A sense related to (25) is ‘to occur regularly/persistently’ in (28), differing from (25) since the motion does not occur in the spatial domain.

- (28) a. *Naturally curly hair runs in my family* [n=5]
 b. *two streams of development run through the history of twentieth-century American folklore*

While many of these senses are related to the central sense(s) via straightforward metaphorical mappings, other senses of *to run* involve yet other patterns of extension. One frequent sense of *to run* can be paraphrased as ‘to function’; the link to the central sense is probably that one can often see that machines are functioning because they and/or their parts move; cf. (29).

- (29) *The monitors ran twenty-four hours each day* [n=47]

A more abstract extension denotes ‘to be valid’ as exemplified in (30).

- (30) *But within that period you must have applied for a new one [a vehicle licence], to run from the day after the last one expired* [n=4]

While I did not discuss all details of how different intransitive senses of *to run* are related to each other, they do appear to form some kind of network like those posited for many other words. We now turn to transitive uses.

3.2. Transitive uses of *to run*

The intransitive uses of *to run* discussed above account for about 65% of all uses in my corpus. Most of the transitive uses we discuss now are related to one of the intransitive uses by what Levin (1993: Section 1.1) refers to as “[o]bject of [t]ransitive = [subject] of [i]ntransitive [a]lternations” as in (31).

- (31) $NP_i V_{\text{intransitive}} \rightarrow NP_k V_{\text{transitive}} NP_i$

For a straightforward motion verb such as *to run*, this is exactly what a cognitively-oriented approach predicts, given the prototypical meaning of the transitive construction (cf. Rice 1987) and the fact that omnipresent conceptual metaphors such as STATES ARE LOCATIONS, CHANGE OF STATE IS CHANGE OF LOCATION and CHANGE OF STATE IS CONTROL OVER AN ENTITY RELATIVE TO A LOCATION can be easily exploited; hence, I will proceed in the order of senses in Section 3.1.

The prototypical sense of ‘fast pedestrian motion’ can also be found in transitive uses where the direct object usually is an event such as a race or a marathon (cf. [32a] and Levin’s [1993: 43-4, 266] locative preposition drop alternation) or a concrete object determining the direction/endpoint of movement (cf. [32b], which can then be paraphrased as ‘to score points by running to some location’), and (32c). Finally, the direct object can also be a distance/measure phrase (cf. Levin [1993: 266]) as in (32d).

- (32) a. *and it’s Sibor who leads the way running his own race* [n=12]
 b. *Or even the way you run bases* [n=1]
 c. *Running the rail in the yellow is Honey Church*
 d. *His brother ran a mile to get the father*

The most straightforward and most productive causative extension from one of the central senses follows from the induced action alternation (cf. Levin [1993: 31]) and is correspondingly instantiated by the senses ‘to cause motion’ in (33) and ‘to knock over’ in (34), which speaks in favor of ‘motion’ as the sense from which most others can be most economically derived.

- (33) a. *He ran a finger down his cheek, tracing the scratch there* [n=13]
 b. *Suppose he ran up the white flag altogether?*

- (34) *The hospitals contain patients [...] run over by sports cars* [n=3]

In the examples in (33), the direct object refers to the thing that is moved, but there are also cases where it is not the direct object that moves; consider (35) for an example of such a profile shift.

- (35) *the soldiers were ordered to [...] run through anyone who might step out of line* [n=1]

In the special case where the direct object is an eye (cf. [36]), *to run* means ‘to see’: directing one’s view to the stimulus is conceptualized as moving the sensory organ to the stimulus which, upon contact, is understood as being perceived via the PERCEIVING IS TOUCHING metaphor.

- (36) *He ran his eye along the roof copings* [n=1]

The sense of ‘to cause motion’ in (33) in combination with the conceptual metaphor MORE IS UP results in the sense ‘to cause to accumulate’ or ‘to increase’ shown in (37); by extension, if the direct object refers to a part/piece of clothing, *to run up* can also mean ‘to sew’ as in (38).

(37) *thanks partly to George Herman Ruth’s spectacular efforts each season to run his salary higher and higher* [n=1]

(38) *Do you love to run up a hem, sew on buttons, [...]?* [n=1]

Other transitive uses constitute causative extensions from metaphorical intransitive uses. For instance, the metaphorically motivated sense of ‘to deteriorate’ has a transitive counterpart ‘to cause to deteriorate’, which is exemplified in (39), and since the mental state of a human being can be worsened by, e.g., criticizing somebody, *to run* has also taken on this sense (cf. [40]).

(39) *Have you had the flu or you’ve been [...] run down in the last few days* [n=4]

(40) *Casey had made a point of running down all such suggestions* [n=2]

There is also a fairly fixed transitive extension of the intransitive sense of ‘diffusion of color’ exemplified in (11b) above, namely a sense where *color* becomes the direct object rather than the subject as above (i.e. another alternation of the type schematically represented in [31]).

(41) *the bright V woven into the neckline had melted, running a darker color* [n=1]

A further example of a causative extension of an intransitive sense of *to run* is exemplified in (42), where it means something like ‘to cause something to have a particular wording’, i.e. ‘to formulate’.

(42) *We usually run a social note when somebody moves away* [n=1]

The same mechanism underlies the extension from ‘to extend spatially’ (cf. [24] above) to that of ‘to cause to extend spatially’ in (43).

- (43) *But anybody who promises a substantial volume of business can get a railroad to run a short spur to his plant these days* [n=1]

Then, a sense that could be related to both that of ‘to function’ and ‘to amount to’ and/or that could be explained with reference to Levin’s locative preposition drop alternation is exemplified in (44).

- (44) *To continue to run a public sector surplus, although [...]* [n=3]

The most frequent group of transitive senses of *to run* are causative extensions of ‘to function’. One can be paraphrased as ‘to execute/operate’, the other as ‘to manage’ (cf. [45] and [46] respectively).

- (45) a. *Very often the screens are run at too high a brightness level which can quickly tire the eyes and wear out the screen* [n=25]
 b. *Presently they had to give up running the furnace at full capacity*
- (46) a. *she often saw him when she was in Ramsford, [...], where he ran the one-man police station* [n=101]
 b. *The club runs regular trips to the cabins*

The difference between the two already emerges from the nearly synonymous paraphrases. On the one hand, the sense of ‘to execute/operate’ usually involves starting some machine or (software) application which can then operate on its own or on the basis of continuous personal/manual involvement of the operator. On the other hand, the sense ‘to manage’ usually involves directing some organization or institution on a more abstract level of involvement. Finally, there are some instances where it is not really possible to decide which degree of involvement and, thus, which of these two senses is instantiated; consider (47) as an example (which supports Hanks’s [2000: Section 7] discussion of semantic indeterminacy).

- (47) *When we are able to run a four day first aid course* [n=23]

It is unclear whether (47) means ‘we taught the course (ourselves)’ (as in, e.g., *I’ve got to run an errand*) or ‘we organized the course and let other people teach it’. If we hypothesize that language users have abstracted away from such vagueness, the hypothesized more schematic sense they have stored could be labeled ‘to be in charge of something’.

A further extension of the sense of ‘to execute/operate’ involves a frame addition (cf. Norvig and Lakoff 1987: 197) of what might be called the publication frame, resulting in the sense of ‘to broadcast/publish’ as in (48).

- (48) *The island's newspaper runs a weekly cartoon showing the adventures of 'Vincey'; in its struggle to survive* [n=5]

Finally, *to run* can be used as a transitive phrasal verb with the particle *off* meaning ‘to copy’. While there is some semantic relation of this sense to that of ‘to execute/operate’, the reason for why the particle *off* is part of this construction remains opaque to me; this is probably motivated by the movements which were once involved in the action of copying.

- (49) *If you give me a tape I've got a tape to tape and I can run it off* [n=1]

3.3. More idiomatic uses of *to run*

This section discusses some senses that, while they can of course also be characterized in terms of transitivity, are semantically much more difficult to integrate into the network, given their lack of compositionality. Since many of these senses are also strongly associated with particular content words as complements,¹¹ do not appear to be very productive syntactically, and describe recurrent situations of social interest, they qualify as idioms (cf. Nunberg, Sag, and Wasow 1994: 492–493), deserving special mention in their own section. One of these is the sense ‘to risk’ as in (50).

- (50) *They were reluctant to appoint sheriffs to protect the property, thus running the risk of creating disturbances* [n=12]

Then, there are several idioms which are used to characterize humans’ experiences. In (51), *to run the gamut* refers to ‘to experience a wide variety of things’; in (52), *to run the gauntlet* means ‘to experience being criticized by (many) people’, and (53) denotes ‘experiencing something very negative’.

- (51) *it [red wine] will have run the gamut of many beguiling and interesting stages* [n=3]

(52) *William and Hamrick did indeed run the ga[u]ntlet* [n=4]

(53) *their cups were already running over without us* [n=1]

A different idiom meaning to ‘to ignore’ is exemplified in (54).

(54) *Catholics run roughshod over Protestant sensibilities, by failure to consider the reasoning behind the Protestant position* [n=1]

The final idiom depending on particular content words can be paraphrased as ‘to be successful’ as in (55).

(55) *New Halen running a blinder up in the third* [n=1]

For a representation summarizing the discussion so far, consider Figure 1. Solid lines denote instance and similarity links, dotted lines denote causation alternation links. Note that Figure 1 serves expository reasons only – it is, just like Bartsch’s (1984: 48) polysemic complex, merely a notational format and is non-committal with respect to issues of mental representation.

4. Case studies

This section will introduce several very brief case studies discussing the interplay between the behavioral profile of *to run* and the cognitively-motivated senses. As mentioned above, I will not restrict my analysis to R1 collocations as Kishner and Gibbs (1996) but will base it on a much wider variety of ID tags. To that end, the data set, all instances of the lemma *to run* discussed in Section 3 above, were coded for the following direct ID tags; cf. Divjak and Gries, to appear, for a similar way of annotation):

- morphological features of the verb form: tense, aspect, and voice;
- the syntactic properties of the clause the verb form occurs in: intransitive vs. transitive vs. complex transitive use of *to run*, declarative vs. interrogative vs. imperative sentence form, main clause vs. subordinate clause (e.g. regular subordinate clause with or without subordinator, relative clause with or without relative pronoun);

- semantic characteristics of the referents of the elements co-occurring with *to run*: its subjects/heads, objects and complements (which were coded, e.g., as human, animate, concrete countable objects, concrete mass nouns, machines, abstract entities, organizations/institutions, locations, quantities, events, processes etc.);
- the instance's collocates in the same clause;
- a paraphrase of *to run*'s meaning in the citation.

As a result, I obtained a corpus-based behavioral profile of *to run* based on 815 citations annotated with respect to 252 different ID tags (40% consisting of manually annotated formal and semantic properties mentioned above, 60% consisting of collocates); since the *absolute* sense frequencies varied considerably, I used the *relative* frequencies of each ID tag attribute within each ID tag. The following brief case studies exemplify how these data can be put to use in order to address a variety of questions that virtually all cognitively-oriented analyses of lexical polysemy must address; these questions include the issue of prototype identification, the (degree of) sense distinctness, the structure of the hypothesized network etc.

4.1. Prototypicality of one sense

Let me begin with the question of which sense of *to run* is the prototypical one. This question plays a central role in cognitive-linguistic analyses so various researchers have established a variety of criteria (cf., e.g., Rice [1996: 145–146], Tyler and Evans [2001: Section 3.3]); the following is a non-exhaustive list of such criteria: asymmetrical judgments of goodness or similarity; ease of elicitation; gradation within the category; earliest attested meaning; centrality/predominance in the semantic network; use in composite forms; etc. However, given such an inventory of criteria, conflicts of criteria are the rule rather than the exception (cf. Corston-Oliver [2001] on *by*). The present subsection illustrate how corpus data can be brought to bear on this issue.

Following the argumentation by Norvig and Lakoff (1987: 198) as well as Tyler and Evans (2001, 2003), Figure 1 suggests that 'motion' is the prototypical sense since 'motion' is the sense from which most others can be (most economically) derived. However, both corpus data in general (i.e. data for which no behavioral profile in Hanks's (1996) sense is necessary) and the behavioral profile of *to run* in particular point nearly uniformly into

a different direction, namely that, as I claimed above, *to run*'s prototypical sense is instantiated by 'fast pedestrian motion' as in (4). Let us begin with some arguments from general corpus data.

First, the sense 'fast pedestrian motion' is the most frequent sense used in early stages of acquisition, as is shown by an as yet informal analysis of *to run* in the Manchester component of the CHILDES corpus (cf. Theakston et al. [2001] and MacWhinney [2000] respectively), where the only other sense coming close to a similar frequency is that of 'to knock over'.¹²

Second, according to etymological dictionaries, which are based on the analysis of historical texts and, thus, adopt a corpus-based approach, the "exact sem[antic] and [phon]ological originations and interactions are at once complicated and obscure" (Partridge 1961: s.v. *run*), but the diachronically primary senses are 'fast pedestrian motion' and 'to flow'.

Third, a related argument is that, like so many other English verbs, *to run* can be zero-derived to function as a noun, a development which apparently began in the 14th or 15th century. There are 60 such instances in the ICE-GB corpus (*run* and *runs* occur 47 and 13 times respectively),¹³ about 75% of which refer to 'fast pedestrian motion' (or the metonymically related sense of 'to score in baseball, cricket etc.');

the few exceptions to this predominance are mainly instances from just one corpus file where *run(s)* refers to experimental trials and a few fixed expressions such as *in the long run*. Also, the 'fast pedestrian motion' sense of the zero-derived noun appears first diachronically (cf. the OED 3.01 on CD-ROM, s.v. *run*).

Nearly all of the general corpus data already point in the same direction, but we can also exploit the behavioral profile for further evidence. For example, the data show that the sense 'fast pedestrian motion' is by far the most frequent one in the corpus (approximately 25% of all instances), which reflects its central status (cf. Durkin and Manning [1989]). Also, it appears to be the formally least constrained sense and can, thus, be considered unmarked and prototypical (cf. especially Lakoff [1987: 60–61] on the relation between prototypicality and markedness).¹⁴ But what does "formally least constrained" mean and how can it be measured? One rather specific example for "formally least constrained" is that the sense 'fast pedestrian motion' is the one with the highest number of differently headed prepositional phrases. A more general, and thus more valuable, finding is that 'fast pedestrian motion' is the sense with the highest number of different 252 ID tag attributes, i.e. it exhibits most variation across all formal and semantic characteristics which were coded, which in turn strongly supports its unmarkedness.

In sum, the above arguments demonstrate the utility of corpus data for prototype identification. While my earlier work on this issue has been concerned with the corpus-based identification of prototypical instances of constructions (verb-particle constructions as well as ditransitives and prepositional datives; cf. Gries [2003b, 2003c]), this work illustrates a similar potential for prototypical intraword senses.

4.2. Distinctiveness of senses

A notoriously problematic issue arising for every polysemy analysis is to decide whether two different citations instantiate distinct senses or just modulations of a more general sense (the lumping vs. splitting issue). Some studies (e.g. Tyler and Evans [2001, 2003]) have addressed this issue by invoking the notion of inferrability by arguing that a particular use of the preposition *over* constitutes a different sense if it profiles a spatial configuration and if the meaning of *over* cannot be inferred from encyclopedic knowledge and contextual information. Let me briefly exemplify this approach on the basis of *to run*. Consider, for example, the meaning of ‘to flow’. Given the high degree of granularity of some cognitive-linguistic analyses, one can assume that any cognitively-oriented polysemy analysis of *to run* adopting the full-specification approach would postulate the existence of this sense. However, as the corpus data reveal and as one would expect intuitively, all of the instances of *to run* meaning ‘to flow’ have as their subject/head a (usually uncountable) noun denoting a liquid. Since the only natural way for liquids to move is by flowing, the sense of ‘to flow’ is inferable, which in turn would obviate the need to posit a separately stored sense of ‘to flow’ – rather, positing the sense ‘motion’ is sufficient since the particular kind of motion is contributed by contextual information tapping into encyclopedic knowledge. A similar line of reasoning applies to the sense of ‘to overflow’ with the liquid as subject as in *The water ran over*. Once the meaning of ‘to flow’ is considered compositional, the sense of *The water ran over* can also be inferred from the meaning ‘to flow’ and the independently established sense of *over* in this expression (cf. Tyler and Evans 2001: 756–757). This argument also ‘explains’ why no one has, on the basis of examples such as *Back with Gary Pallister who just let the ball run across the touch-line*, ever postulated that *to run* has a sense ‘to roll’ – the manner of motion is again contingent on (the nature of) the subject.¹⁵

In spite of the intuitive appeal of this approach, it is worth pointing out that this approach appears to run into problems once it is combined with Tyler and Evans's approach to prototypicality: on the one hand, they should consider the sense 'to flow' prototypical (since it is among the earliest attested senses), but on the other hand it should not be considered an individual sense in the first place (since it is inferable); the latter position, however, I would contend, does probably not really match native speakers' intuitions.

From the corpus-based perspective, however, such a conflict does not even arise in the first place. In a paper on linguists' contribution to questions of how polysemous words are mentally represented, Croft (1998: 169) refers to the senses of *to eat* labeled 'to consume' and 'to dine', arguing that the comitative argument (referring to a fellow eater) occurs only with the latter use in the corpus. That is, one finds sentences of the type *Jack ate lunch with Jill* but not *Jack ate a pizza with Jill*, although the latter would be judged grammatical on introspection. The disjoint syntactic-semantic distribution suggests that 'consume' and 'dine' are grammatically distinct uses of *eat*. (Note how this argument obviously presupposes some version of a behavioral profile of *to eat*.)

If we apply this argument to the example of 'to flow', the citations involving the sense 'to flow' all have subjects being a liquid so that, in accordance with intuitions, this distributional characteristic provides corpus-linguistic evidence for considering 'to flow' a distinct sense. In spite of some lexicographical implications,¹⁶ this example is relatively trivial: even if the subject of *to run* meaning 'to flow' always has a liquid in the behavioral profile, one does not need the behavioral profile to find that out. However, there are less trivial examples to drive home the point that corpus data help to distinguish senses in terms of formal patterns so let us now look at two such examples, one in favor of lumping, one in favor of splitting.

First, this approach would argue in favor of lumping two kinds of usage of 'fast pedestrian motion.' For example, we find cases where *to run* in its sense of 'fast pedestrian motion' is combined either only with a SOURCE argument (cf. [56] and n. 13) and or only with a GOAL argument (cf. [57]).

(56) *and we ran back to my car*

(57) *Durkin and Calhoun came running from the post.*

But, on the basis of the above arguments, we must not infer that the two constitute two different senses of *to run* since there are many examples such as those in (58), where both SOURCE and GOAL are present (in both orders), and, in fact, such cases outnumber those with only a SOURCE.¹⁷

- (58) a. *He was almost breathless from having run towards her uphill from, it could only be, the lake*
 b. *I once ran from the Archive studio to the Start The Week studio*

Second, this approach would argue in favor of splitting the two senses of *to run* that can be paraphrased as ‘to move away from something dangerous/unpleasant’ and ‘to move away to engage in a romantic relationship’. The data show that the former sense is instantiated by the verbs *to run off* and *to run away* (often with a prepositional phrase referring to the negatively evaluated stimulus). The latter sense also occurs as *to run off* and *to run away*, but mostly with a comitative argument. The parallel to Croft’s example is that, while a sentence with both a negative stimulus and a comitative argument (e.g. *She ran away with him from all the problems*) appears acceptable, not a single such sentence was attested in my data, which, following Croft’s logic, points to the distinctness of the two senses. Again, objective corpus-based evidence could be used to answer an otherwise difficult question or, more modestly, could provide objective *prima facie* evidence in one direction.

4.3. Where to connect a sense in the network

Another interesting possibility of analysis arising from the behavioral profile is concerned with determining the structure of the network representing the senses of *to run* and their relations. Consider again the senses ‘to move away from something dangerous/unpleasant’ and ‘to move away to engage in a romantic relationship’. Devising a lexical network structure of *to run* requires a decision how to connect these two senses to the others. The initial decision would probably be to connect them to the node of the prototypical sense ‘fast pedestrian motion’ since this is the central sense and ‘fast pedestrian motion’ is the typical/most basic way to perform these actions. On the other hand, it is equally obvious that fast pedestrian motion is not the only way to move away from something dangerous/unpleasant or to move away to engage in a romantic relationship, which is why the senses ‘fast motion’ (or just ‘motion’) appear reasonable points of connection, too.

It is therefore difficult to decide in favor of one of the two alternatives on a principled basis, i.e., to decide how to integrate them into the network such that they connect to the sense they are most similar to. However, one can approximate the semantic similarity of these five senses in terms of their distributional similarity (as is customary in corpus-linguistic or computational-linguistic studies; cf. Biber [1993] for an example and McDonald [1997] for validation).

Previous studies aiming at quantifying the similarity of senses have used hierarchical cluster analyses on semantic similarity judgments or sentence sorting tasks. For the moment, however, the simpler technique of correlation analysis also serves our purpose. I computed all 3,080 pairwise correlations of the 56 senses' ID tag vectors to determine whether this approach is feasible at all. The results support this (in other areas already well-established) approach in many respects: First, the correlation coefficients obtained range from .38 to .93, differentiating across a whole spectrum of degrees of distributional similarity. Second, a brief look at the extreme values shows that the senses least similar to each other are those in (59a) and (59b), an intuitively reasonable result.

- (59) a. *their cups were already running over without us*
 b. *He ran his eye along the roof copings*

Third, the result concerning the senses considered most similar to each other by this correlational analysis appears to be even more reasonable, and it also bears directly on our question: the maximum r value (i.e. the highest degree of similarity) is obtained for 'fast pedestrian motion' and 'to escape'. Finally, the five senses we are interested in are on average much more similar to each other than the average pairwise similarity of senses after Fisher Z transformation (mean $r_{\text{all senses}} = .545$; mean $r_{\text{five senses}} = .848$), as would again be expected intuitively. These results lend credence to the assumption underlying much recent corpus-linguistic work that distributional similarity correlates with semantic similarity.

However, the results are rendered less precise than possible for our actual question since many ID tags occur so infrequently that their percentages are by definition either very small or very large, thereby distorting the results. Thus, I left out the ID tags coding just the presence/absence of a particular adverb or preposition, which left 55 reasonably frequent ID tags for comparison. Then, I computed the correlations between the three motion senses and the two 'to escape' senses (across all 55 relative frequencies).

The results are unequivocal: across all 55 ID tags, the sense ‘to move away from something dangerous/unpleasant’ is highly significantly more similar to ‘fast pedestrian motion’ than to the senses ‘fast motion’ ($z=5.38$; $p<.001$) and ‘motion’ ($z=5.06$; $p<.001$) while the latter two do not differ significantly from each other ($z=.45$; $p=.665$). The same holds for the sense ‘to move away to engage in a romantic relationship’, which is very significantly and marginally significantly more similar to ‘fast pedestrian motion’ than to the senses ‘fast motion’ ($z=3.17$; $p=.002$) and ‘motion’ ($z=1.88$; $p<.061$) respectively while, again, the latter two do not differ significantly ($z=1.42$; $p=.156$). That is to say, in absence of further theoretical motivation or evidence to the contrary, one should connect both ‘to escape’ senses to the prototypical sense rather than to ‘motion’ or ‘fast motion’, a decision we could again motivate on the basis of objective evidence.¹⁸

4.4. Agglomerative clustering of senses

Cluster analyses have been used to determine the similarity of intraword senses or the degree of granularity exhibited by polysemous word senses (cf. Miller 1971; Sandra and Rice 1995; Rice 1996). While clustering is often applied to collocate frequencies (cf. Manning and Schütze 2000: ch. 14), we can apply it to the much more detailed complete behavioral profile of *to run*; cf. Schulte im Walde (2003) for a similar approach. Accordingly, the table of relative frequencies was submitted to a hierarchical agglomerative cluster analysis, resulting in the dendrogram in Figure 2.¹⁹

Given the limited corpus size, the results can only be preliminary, but in spite of the diversity of authentic corpus data, several noteworthy observations can be made. First, on the right we find a branching which corresponds extremely closely to that of intransitive and transitive (i.e. causative) uses.²⁰ Then, at the top of Figure 2, the analysis has grouped together most cases of literal motion and a range of cases of abstract motion that have in Section 3 been related straightforwardly via metaphorical mappings and/or image-schema transformations.

Within this larger cluster, several small ones are homogeneous enough to be mentioned: ‘fast pedestrian motion’ and ‘to escape’, ‘to extend spatially’ and ‘motion’, and ‘motion without control/restraint’ and ‘metaphorical motion without control/restraint’. There is also a cluster subsuming several semantically very similar senses under some general ‘to be in charge of’ sense. Finally, there is one cluster subsuming four senses which




Figure 2. Dendrogram resulting from a hierarchical cluster analysis

all denote that some harm is done to the referent of the direct object; this is of course still a heterogeneous class, but the commonality is striking. In other words, we do find several clusters at about the same level of generality as Sandra and Rice (1995) do in their work on prepositions.

Other findings from Figure 2 are also worth mentioning. For example, the senses amalgamated earliest on distributional grounds are exactly those that are most strongly branching in the network-like representation of Figure 1. Note that this cannot be explained in terms of the senses' frequencies and that particularly the sense of 'fast pedestrian motion' argued to be prototypical is the very first sense to be amalgamated. If this finding could be replicated, it would open up completely new perspectives since corpus-based clustering of objective ID tags can then be considered indicative of, or again more modestly at least correlating with, aspects of category structure.

I hope the brief comments above have outlined the potential of the method. However, given the limited size of the present corpus, the larger data set necessary for a more comprehensive analysis may well result in some changes. For example, there are several high intercorrelations (of the type discussed in Section 4.3) and some clusters which are difficult to explain (e.g., the cluster linking "to increase" and "to broadcast", which are, however, not amalgamated early). Part of such variance is of course due to the fact that cluster analyses are influenced by the amount of noise attributable to corpus data. But since research on clustering of humanly-sorted word senses has also sometimes resulted in low agreement ratios (cf. Jorgensen 1990), the corpus-based method is not by definition an inferior method; currently ongoing work tests corpus-based clusterings of the above sort against experimentally obtained sorting preferences. Be that as it may, depending on the size of the data set and further ID tags one might wish to include (e.g. metaphorical mappings or other mechanisms underlying sense extensions), future analyses can shed much more light on categorical and distributional properties of particular senses.

4.5. Automatic sense identification

I repeatedly referred to the fact that Kishner and Gibbs have argued in favor of adding lexico-grammatical information to the description of (the interrelations of) polysemous words' senses. While such a compilation of a behavioral profile would obviously not only enhance the descriptive adequacy

of the analysis as such, the previous sections have, I hope, also demonstrated that the behavioral profile offers a variety of possibilities to arrive at objective answers to notoriously difficult questions. However, such an approach has more to offer once we begin to leave the domain of cognitive linguistics proper. In other words, while we have so far considered the senses as given and have then determined ID tags differentiating between senses, we can also adopt the reverse perspective: how well can we predict the sense of *to run* in a particular citation when we extract this citation's ID tags? More technically, so far we looked at the conditional probability $p(\text{ID tag} \mid \text{word sense})$, but we can equally well determine $p(\text{word sense} \mid \text{ID tag[s]})$, a question central to the issue of word sense disambiguation (WSD) within computational linguistics. If the joint predictive power of several ID tags made it possible to predict a word sense, this would provide further support for the notion of ID tags and analyses relying on them. This is how a very elementary approach would look like; cf. Manning and Schütze (2000: Chapter 7) for much discussion of WSD.

The regular way to determine intercorrelations between a particular meaning of *to run* on the one hand and formal/semantic patterns of the sentences instantiating this meaning on the other hand would be to cross-tabulate all meanings with all ID tags. However, given the large number of potentially relevant features, the number of possible configurations (of different factorial degrees) of ID tags increases so quickly that the observed frequencies for each configuration turn out to be too low to lend themselves to usual statistical approaches (e.g. the χ^2 -test), a frequent problem in such applications. In order to overcome a similar problem, Gildea and Jurafsky (2001: Section 4.2) suggested to combine probabilities of a selected variety of meaning-pattern configurations, a technique we can adopt easily.

Assume that a sense recognition system is provided with (i) a general baseline frequency of each sense and (ii) a mechanism to identify ID tags of each sense on the basis of the context of the word. When the system is fed with a sentence to recognize its sense, two things can happen. The usual case would be that the ID tag is not particularly distinctive for, i.e. independent of, the sense, and thus just adds noise to the classificatory problem since combining independent probabilities requires their multiplication. The interesting cases are those where the ID tags recognized by the system are not independent of the sense and, in spite of the statistical tendency of the probabilities to decrease, actually *increase* the system's predictive power.

Let me explain this briefly on the basis of the two most frequent senses of *to run*. The most frequent sense in the present corpus (25% of all tokens) is that of ‘fast pedestrian motion’. An automatic sense classification system could already achieve an accuracy of about 25% by simply assigning this sense to every incoming sentence with *to run*. But ‘fast pedestrian motion’ has some strong probabilistic ID tags: If the system recognizes that the verb is in the past tense (i.e. *ran*), then, since 42.3% of all occurrences of *ran* are instances of ‘fast pedestrian motion’, the prediction accuracy rises to 42.3%, an improvement of approximately 70%. If the system also recognizes that *ran* is used intransitively, the prediction accuracy is further increased to 49.3%, and if intransitive *ran* is followed by a prepositional phrase headed by *to*, the prediction accuracy is increased to 73.7%. Finally, if the structure [S [NP_{subj}] [VP *ran* [PP *to* [NP]]]] has a human subject noun phrase, the only attested sense for this configuration of ID tags is in fact ‘fast pedestrian motion’, i.e. the prediction accuracy amounts to 100%. A similar case can be made for the second most frequent meaning of *to run*, ‘to manage’. Its overall relative frequency is 12.4%, but, as is shown in Figure 3, there are several formal, easy-to-recognize ID tags strongly associated with the meaning of ‘to manage’.

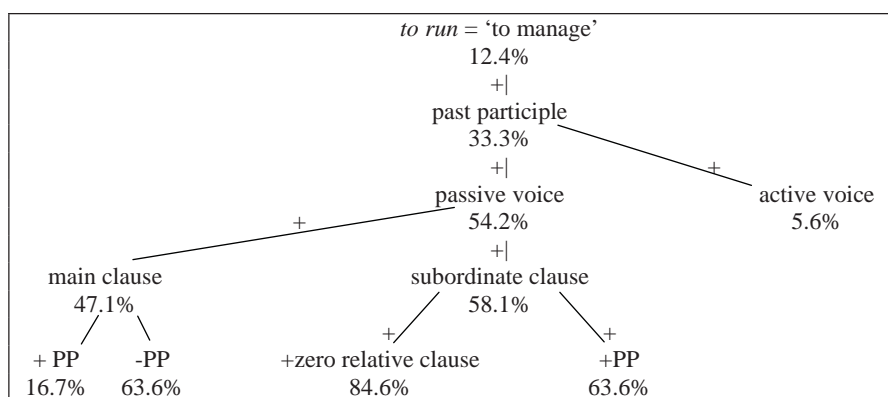


Figure 3. Successive change of prediction accuracy for ‘to manage’

Space does not permit detailed discussion of more examples; suffice it to say that similar accuracy improvements are obtained for other sufficiently frequent senses (such as for ‘to extend spatially’). In addition, while the above observations have exclusively relied on positive ID tags, i.e. ID tags whose *presence* is probabilistically indicative for a particular sense, senses

can of course be equally well predicted given the *absence* of some ID tag; for reasons of space, I will not discuss this phenomenon in detail, but cf. Table 1 for results on the two most frequent senses.

Table 1. The complementary distribution of selected ID tags for the two most frequent senses of *to run*

		'fast pedestrian motion'	'to manage'
verb form	+	<i>ran</i>	<i>run</i> (past part.)
	-	<i>runs</i>	<i>ran</i>
transitivity	+	intransitive	transitive
	-	transitive	intransitive
clause type	+	main clause, imperative clause	zero relative clause, zero subordinate clause
	-	(zero) relative clause interrogative clause	main clause
subject	+	human, animate	organization/institution
	-	concrete objects, organization/institution	-
preposition of following PP	+	towards, for, down, after, up	0
	-	-	-

These examples show that the senses of *to run* do have strong probabilistic associations with formal and/or semantic patterns (in our parlance, ID tags) that, according to Kishner and Gibbs, merit inclusion in polysemy networks (if only for the sake of completeness and the utility they may have for explaining psycholinguistic findings and their predictive power for NLP); cf. also Theakston et al. (2002) for similar correlations concerning the acquisition of the verb *to go* as well as Newman and Rice (this volume) for such findings concerning the verbs *to eat* and *to drink*. The most interesting thing about Table 1, however, is that the ID tags also differ across senses unexpectedly: there is no *a priori* expectation that 'fast pedestrian motion' should correlate with past tense whereas 'to manage' and 'to extend spatially' correlate with past participle and third person singular present tense/the present participle respectively. True, it is easy to motivate that 'to manage' is the only sense associated with past participle in passives: it is the only transitive, and thus passivizable, sense of the two singled out for analysis, which is in turn responsible for the expected transitivity preferences. But note that (i) there would nevertheless not have been a reason to posit that 'to manage' is associated with past participle to begin with – could it not equally well prefer the present participle? – and (ii) it does not

explain other senses' ID tags or their complementary distributions. In sum, the mere fact that ID tags are powerful enough for the actual prediction of senses already underscores their importance for the analysis of word senses.

5. Conclusion

Before I summarize the most important points of this paper and briefly talk about possible extensions, one caveat is necessary. I have tried to emphasize the benefits of additional corpus-based evidence, but I should like to point out, however, that I do not advocate using corpus evidence alone. Corpus evidence can complement different research methodologies such as (psycho-)linguistic experiments, but it should not replace them. Thus, not all results from above will remain constant across different data, and the above findings will have to be checked against different evidence.

There is yet another aspect of the corpus-based approach that deserves mention. Given the multifactorial approach advocated above, I am the first to admit that the present corpus is not large enough, which is also reflected by the fact that not all senses of *to run* listed in reference works were attested. But in spite of this, it was large enough for us to find (i) which senses are most frequent and whose characterization and cognitive motivation is therefore most relevant and (ii) that some uses of *to run* attested in the corpus data were not listed in (corpus-based) dictionaries such as Cobuild on CD-ROM (1995) or Collins Cobuild E-Dict (1998) and/or turned out to be unfamiliar to some native speakers (e.g. the 'to stop talking' sense and the 'to fail to understand' sense). This is all the more astonishing since these senses or (some aspects of) their distributional behavior are not fully predictable from other senses and should thus be listed in reference works. I interpret this as evidence that it is highly unlikely that intuitions of linguists concerning (i) what are possible uses of a lexeme and (ii) how frequent (or, more cognitively speaking, how entrenched) the uses are will turn out to provide a data base reliable enough for analyzing a word's senses.

I hope the brief case studies discussed in the previous sections have borne out my claim that cognitively-oriented analyses of polysemy benefit from a corpus-based perspective. The hypothesis by Kishner and Gibbs (1996) that comprehenders' choices of a sense of a polysemous item can be influenced by senses' colligations has received support. Also, we have seen that there are many recurrent problems of polysemy analyses (which sense

is prototypical?, where do we connect sense X to the network? etc.) to which corpus-based methods can contribute their share of an answer. What is more, I hope to have shown (i) that Kishner and Gibbs's proposal to provide the lexical network description of words' senses with ID tags and frequency information is in fact reasonable and (ii) how this proposal could be implemented since it could be demonstrated that senses often have so strong associations to ID tags that (combinations of) ID tags alone suffice for automatic word sense identification.

Finally, the above observations underscore individual word senses' strong affinity to constructions. The present approach, therefore, bears quite some resemblance to the analyses of collocation strength (cf. Stefanowitsch and Gries [2003] and Gries and Stefanowitsch [2004a, 2004b]). In these studies, the focus was on measuring the degree of attraction and repulsion of words and constructions by determining the words most characteristic for particular constructions. However, they also point out that the reverse perspective – starting out from a word's behavioral profile – is equally possible. The present study follows up on that proposal: It starts from a single word, and it measures the degree of attraction/repulsion of this word's senses and ID tags, thereby taking into consideration, and simultaneously supporting, recent findings indicating that some distributional patterns are often not *verb*-specific but rather *verb-sense* specific (cf. Roland and Jurafsky [1998, 2002] and Hare, McRae and Elman [2003]). Since these studies are, however, mostly concerned with words having few senses which are much less similar to each other than those of *to run*, the present work with its much more comprehensive and cognitively-oriented behavioral profile also contributes its share to the large area of disambiguation preferences in language comprehension as discussed by Roland and Jurafsky (2002) and Hare, McRae and Elman (2003: 282–284, 295–298). Some possible extensions of this approach will be proposed below.

Let me begin with a major methodological suggestion for improvement. The main multifactorial technique employed above has been the hierarchical agglomerative clustering technique; its main emphasis has been on determining degrees of similarity between (groups of) senses. In passing, we have also briefly looked at to what degree senses can be predicted on the basis of ID tags. Since the number of variables strongly correlating with meanings of *to run* has been very high, it was – as in Gildea and Jurafsky's work – not possible to include all possible combinations of features in the analysis for the latter objective. Therefore, I isolated some ID tags with a strong predictive power on the basis of a manual inspection of hundreds of

two- and three-dimensional frequency tables. Given the large number of tables of different factorial degrees requiring manual inspection, this technique must be refined for objective analysis. On the basis of a much larger data set, comprehensive statistical techniques widely used in other disciplines could be used, which I would like to briefly outline in the following.

A first such technique is known as hierarchical configural frequency analysis (H-CFA). The basic idea of a regular configural frequency analysis (CFA) is similar to that of a χ^2 -test (cf. von Eye [1990] or Krauth [1993] for details). For the present purposes, it determines which of the observed configurations of ID tags of the n-dimensional table are significantly more/less frequently than expected using multiple post hoc tests. For our purpose, however, an extension of this technique, the H-CFA, is more promising: it tests all configurations of ID tags of all factorial degrees. In our case, the analysis would test (i) which of the thousands of configurations of ID tags and senses are frequent enough to be statistically significant and (ii) which of the ID tags are necessary to constitute a particular significant configuration and which can be omitted because they do not discriminate well enough between senses. For example, instances of *to run* meaning 'fast pedestrian motion' are often used intransitively with a human agent – but a temporal prepositional phrase will probably not be distinctive for this meaning. Put differently, human agents and intransitive usage rule out many meanings of *to run* other than 'fast pedestrian motion', but a temporal prepositional phrase denoting the time when the action takes place can be predicated of most meanings of *to run* and is, thus, not useful for the identification of 'fast pedestrian motion'.

Another possible technique which would not require such large sample sizes would not use significance testing but would otherwise generate similar results, namely the technique of association rules frequently used in data mining. Its measures (coverage, support, strength, lift, and leverage) could also identify recurring configurations at different factorial levels.

A much less technical way of extension can be introduced on the basis of two issues already previously mentioned. First, we have seen that recent research on word senses summarized above suggests that it is probably more rewarding to abandon traditional word senses on behalf of meaning components. Second, a more explicit cognitive analysis of *to run* could provide more evidence of the frequencies of mechanisms which figure in extensions of words (i.e. metaphorical, metonymical or image-schematic mappings, profile shifts, frame additions etc.) than the relatively coarse-grained analysis in Section 3. These cognitive mechanisms can be inter-

preted as constituting, or at least contributing to, the meaning components determining a word's sense, and it would be natural in many cases to expect that the derivation of a new sense via some of these mechanisms would manifest itself not just in the abstract analysis of the linguist, but also in one contextual property, which we have labeled ID tags: obviously, many of the abstract motion senses discussed above differ from the literal motion cases such that their subject is not human or animate, to give just one example. On that basis, the kinds of metaphorical mappings figuring in the senses of *to run* would strongly increase (i) the predictive power of sense recognition and (ii) the descriptive power of the clustering algorithm since, then, senses which are metaphorically closely related but otherwise distributionally dissimilar but would receive higher similarity ratings (e.g. 'to function' and 'to be valid'). Probably more interestingly, it would even be conceivable that further studies could investigate which metaphorical mappings are most entrenched and/or exploited most frequently for extending senses and why.

In that connection, these issues may also be interesting from the perspective of language acquisition. For example, setting up a dynamic behavioral profile of a verb while it is acquired may provide interesting evidence concerning both the salience of the different senses as well as the salience of the cognitively-motivated mechanisms underlying sense extensions during acquisition. For example, such a dynamic behavioral profile would provide information as to which senses are acquired first, which metaphors, metonymies etc. are responsible for the first extensions, what the sequence of acquisition of verb senses tells us about the way children extend senses, and how different formal aspects of the words under investigation (e.g. TAM marking, the distribution of senses within different kinds of clauses etc.) figure in their acquisition; the above-mentioned study of Theakston et al. (2002) is a study in this spirit. The same may hold (though perhaps less directly so) for a diachronic approach towards how the different senses of words develop. Finally, while the present approach has focused on behavioral profiles of different senses of the same word, it can also be applied to the corpus-based cognitive-linguistic investigation of near synonymous words (cf. Divjak [2004, this volume] as well as Divjak and Gries [to appear]).

In sum, for many of these issues which are relevant to cognitive-linguistic approaches, a behavioral profile is, I believe, the most rewarding starting point that will hopefully be utilized more fully in future work.

Notes

- * I thank (in alphabetical order) particularly Ewa Dabrowska and Stefanie Wulff, but also the reviewers, for their detailed feedback and comments; the usual disclaimers apply.
- 1. In this paper, I will only focus on the synchronic relatedness of senses.
- 2. The first criterion of this methodology leaves open the question of how word classes other than (usually spatial) prepositions should be investigated.
- 3. This window is much larger than that for most collocation-based studies, but can be justified given the importance of topical context for word sense disambiguation (cf., e.g., Chodorow, Leacock and Miller 2000).
- 4. The ideas of ID tags making up a behavioral profile is also compatible with many studies on word sense disambiguation (cf. Ide and Véronis 1998). Given the overlap of issues that both cognitive linguists and WSD studies have been addressing for quite some time (e.g. distinctness of senses, granularity of sense distinctions etc.), it is even a little surprising to note how little cognitive linguists appear to have looked at the accomplishments of these disciplines. Parts of the present paper will therefore also attempt to bridge this gap.
- 5. This list of frequencies of intraword senses is more than just an end in itself. For example, Williams (1992: 208) discusses “an asymmetry in the amount of priming between central and non-central meanings”, proposing that “[o]ne explanation of this asymmetry would be in terms of the relative frequencies of the two meanings”. However, he is forced to acknowledge that “[i]n the absence of any data on actual frequencies of use of the meanings of these words, this hypothesis cannot be evaluated”. Thus, the knowledge of sense frequencies resulting from the behavioral profile, although of limited use for a traditional cognitive-linguistic analysis of word meaning, are in fact very useful to explain such psycholinguistic findings.
- 6. The dictionaries used were Cobuild on CD-ROM (1995), Collins Cobuild E-Dict (1998) and Merriam Webster’s online dictionary at <http://www.m-w.com>.
- 7. I have included all citations of *to run* into the analysis even if this included very creative uses or complex-transitive uses in verb-particle constructions etc. To my mind, this does not constitute a weakness of the present approach: on the one hand, the importance attached to such less central cases can be weighted by their frequency; on the other hand, a truly cognitively-inspired analysis should be able to provide some motivation for extraordinarily creative or more idiomatic extensions anyway. For an earlier analysis of *to run*’s senses, which invokes prototypes, metaphor, and metonymy from a formal semantics perspective, but addresses only a limited number of senses, cf. Bartsch (1984).

8. If the intention to meet the referent of the prepositional phrase headed by *into* is absent, this sense is probably not only related to the sense of ‘motion’, but probably also to that of ‘motion without control/restraint’.
9. The sense ‘to overflow’ with a liquid as subject (rather than the container) was not attested in the corpus data.
10. Contrary to Lehrer (1990: 226), this sense is neither restricted to, nor significantly preferred in, the corpus of American English.
11. Interestingly, most senses of *to run* that are idiomatic and/or that are tied to particular open-class lexical items involve alliteration: *to run rampant*, *to run riot*, *to run roughshod over NP/S*, *to run the risk*, *to run into rapture*.
12. The analysis counted only complete and completely intelligible utterances by children which were not labeled as imitations or routines; cases where an unambiguous identification was not possible were discarded.
13. The figures result from discounting run/runs as part of proper names – if these were included, the figures would not change markedly, especially since some cases of *run* as part of a proper name refer to racing horses whose main purpose is of course running in the sense of ‘fast pedestrian motion’, thereby supporting my above claim.
14. Note also that (4) is not only prototypical for *to run* because it exemplifies the sense ‘fast pedestrian motion’ – it is also a prototypical instance of the sense of ‘fast pedestrian motion’ because it contains a locative prepositional phrase denoting the goal of the motion like most corpus examples of this sense; corpus examples containing no such prepositional phrase or containing a prepositional phrase denoting the source, direction or origin or goal of the agent’s movement are markedly less frequent.
15. Of course, even though the sense of, say, ‘to flow’ is inferable and need not be stored, it may still be stored just because it is frequent enough to acquire unit status at some point of time.
16. Especially dictionaries often “violate” the criterion of inferrability to provide maximally explicit assistance. I cannot discuss here individual dictionaries’ shortcomings or investigate if dictionaries should prefer listing cognitive mechanisms relating different senses over many minimally different senses, but let me provide just a few examples of debatable decisions in favor of splitting from the Cobuild on CD-ROM (1995). Once the ‘motion’ and ‘to cause motion’ senses of *to run* have been established, do we really need to distinguish the definitions of (a) and (b) from each other, and the senses in (a) and (b) from the one in (c) in spite of their compositionality?
 - (a) ‘motion’: “If an object such as a ball runs somewhere, it moves smoothly and quickly over the ground. EXAMPLE: *The ball ran to the boundary*” (sense 22) vs. “If you run somewhere in a car, you make a short trip in it. EXAMPLE: *Why don’t we run down to Worcester for the afternoon?*” (sense 18) vs., as just discussed above in the main text, “If a

liquid runs somewhere, it flows in a particular place or direction. EXAMPLE: *Tears were running down the side of his face ... The water ran into a bucket.*" (sense 23)

- (b) 'to cause motion': "If you run an object or your hand over something, you make the object or your hand touch it and move over it. EXAMPLE: *He ran his hand over her hair ... She ran her finger down a list of names*" (sense 10) vs. "If you run someone somewhere in a car, you drive them there. EXAMPLE: *Would you mind running me to the station?*" (sense 19)
 - (c) 'motion' or 'to cause motion' (ergative verb): "If you run a vehicle somewhere or it runs there, it moves to a particular place or in a particular direction. EXAMPLE: *Run the car into the garage before you go ... The cart ran down the road out of control.*" (sense 20)
17. In fact, a similar logic can be applied to the potential distinction of *to run* with a DIRECTION argument, which can occur alone (e.g. *He'd heard the shouts and shrieks, had heard Cassie running up the stairs*), but also with a GOAL (e.g. *Russ ran up the steps quickly to the plank porch*).
 18. There is nothing in the method implying that there is just one correct way of analysis or connection: of course, multiple connections between different senses are possible – even then, the proposed way of analysis makes it possible to rank the potential connection sites in terms of similarity.
 19. Hierarchical agglomerative cluster analysis is a family of methods that aims at identifying and representing (dis)similarity relations between different items (a general comprehensive discussion of clustering can be found in Kaufman and Rousseeuw [1990]). Usually, clustering is performed on the basis of variables that characterize the items or on the basis of a similarity matrix of the items as, for example, obtained from the variables in example above or from similarity judgments or sorting tasks. In our case, the items correspond to the senses of *to run* while the variables are the ID tags.

A cluster analysis of the kind used here begins by considering each of the *n* senses as one-sense clusters and proceeds to amalgamate those clusters which exhibit the highest intra-cluster similarity and the lowest inter-cluster similarity successively until all clusters have been amalgamated into a single cluster containing all items. The structure yielded by this amalgamation process is typically represented by means of a so-called dendrogram, i.e. a tree diagram representing the similarities among clusters. In addition, a variety of statistical measures can be outputted that help (i) to determine the number of clusters one should assume as well as (ii) to identify which of the variables are most responsible for the clustering solution obtained.

Since the choice of the distance measure and the clustering algorithm can bias the results, 48 different cluster analyses were conducted to systematically compare different combinations of distance measures, clustering algorithms

and senses and ID tags to include. However, it turned out that in this case the differences were of relatively minor importance. The solution presented above is based on Euclidean distances, the weighted pair-group average, all senses minus the idiomatic ones, and all features.

20. Interestingly, the fact that the coarsest distinction in the corpus data is the syntactic one of transitivity ties in perfectly with a result from a sorting experiment in Miller (1971: 577), where “adult judges seem to work by sorting the items on syntactic grounds before sorting them on semantic grounds.”

References

- Atkins, Beryl T. Sue
 1987 Semantic ID tags: corpus evidence for dictionary senses. Proceedings of the Third Annual Conference of the UW Centre for the New Oxford English Dictionary, 17–36.
- Bartsch, Renate
 1988 The structure of word meanings: Polysemy, metaphor, metonymy. In: Fred Landman and Frank Veltman (eds.), *Varieties of Formal Semantics*, 27–54. Dordrecht: Foris.
- Biber, Douglas
 1993 Co-occurrence patterns among collocations: A tool for corpus-based lexical knowledge acquisition. *Computational Linguistics* 19 (3): 531–538.
- Brugman, Claudia
 1984 The *very* idea: a case study in polysemy and cross-lexical generalizations. *Papers from the Parasession on Lexical Semantics*, 21–38. Chicago, IL: CLS.
- Brugman, Claudia and George Lakoff
 1988 Cognitive topology and lexical networks. In: Steven L. Small, Garrison W. Cottrell, and Michael K. Tanenhaus (eds.), *Lexical Ambiguity Resolution*, 477–508. San Mateo, CA: Morgan Kaufman.
- Chodorow, Martin, Claudia Leacock, and George A. Miller
 2000 A topical/local classifier for word sense identification. *Computers and the Humanities* 34 (1/2): 115–120.
- Cobuild on CD-ROM 1.2.
 1995 Glasgow: Harper Collins.
- Collins Cobuild E-Dict.
 1998 Glasgow: Harper Collins.

- Corston-Oliver, Monica
 2001 Central meanings of polysemous prepositions: Challenging the assumptions. Poster presented at the ICLC 2001, University of California, Santa Barbara.
- Croft, William
 1998 Linguistic evidence and mental representations. *Cognitive Linguistics* 9 (2): 151–173.
- Cruse, D. Alan
 1986 *Lexical Semantics*. Cambridge: Cambridge University Press.
- Divjak, Dagmar
 2004 Degrees of verb integration. Conceptualizing and categorizing events in Russian. Doctoral dissertation, Katholieke Universiteit Leuven.
 this vol. Ways of intending: Delineating and structuring near-synonyms.
- Divjak, Dagmar and Stefan Th. Gries
 to appear Ways of trying in Russian: clustering and comparing behavioral profiles. *Corpus Linguistics and Linguistic Theory* 2 (1).
- Durkin, Kevin and Jocelyn Manning
 1989 Polysemy and the subjective lexicon: Semantic relatedness and the salience of intraword senses. *Journal of Psycholinguistic Research* 18 (6): 577–612.
- Eye, Alexander von
 1990 *Introduction to Configural Frequency Analysis: The Search for Types and Antitypes in Cross-classifications*. Cambridge: Cambridge University Press.
- Fillmore, Charles J. and Beryl T. Sue Atkins
 2000 Describing polysemy: the case of ‘crawl’. In: Yael Ravin and Claudia Leacock (eds.), *Polysemy: Theoretical and Computational Approaches*, 91–110. Oxford: Oxford University Press.
- Geeraerts, Dirk
 1993 Vagueness’s puzzles, polysemy’s vagaries. *Cognitive Linguistics* 4 (3): 223–272.
- Gibbs, Raymond W. Jr. and Teenie Matlock
 2001 Psycholinguistic perspectives on polysemy. In: Hubert Cuyckens and Britta Zawada (eds.), *Polysemy in Cognitive Linguistics*, 213–239. Amsterdam/Philadelphia: John Benjamins.
- Gildea, Daniel and Daniel Jurafsky
 2001 Automatic labeling of semantic roles. *Computational Linguistics* 28 (3): 245–288.
- Gries, Stefan Th.
 2001 A corpus-linguistic analysis of *-ic* and *-ical* adjectives. *ICAME Journal* 25: 65–108.

- 2003a Testing the sub-test: A collocational-overlap analysis of English *-ic* and *-ical* adjectives. *International Journal of Corpus Linguistics* 8 (1): 31–61.
- 2003b *Multifactorial Analysis in Corpus Linguistics: The Case of Particle Placement*. London/New York: Continuum.
- 2003c Towards a corpus-based identification of prototypical instances of constructions. *Annual Review of Cognitive Linguistics* 1: 1–27.
- Gries, Stefan Th. and Anatol Stefanowitsch
- 2004a Extending collocation analysis: A corpus-based perspective on ‘alternations’. *International Journal of Corpus Linguistics* 9 (1): 97–129.
- 2004b Co-varying collexemes in the *into*-causative. In: Michel Achard and Suzanne Kemmer (eds.), *Language, Culture, and Mind*, 225–236. Stanford, CA: CSLI.
- Hanks, Patrick
- 1996 Contextual dependency and lexical sets. *International Journal of Corpus Linguistics* 1 (1): 75–98.
- 2000 Do word meanings exist? *Computers and the Humanities* 34 (1/2): 205–215.
- Hare, Mary L., Ken McRae, and Jeffrey L. Elman
- 2003 Sense and structure: Meaning as a determinant of verb subcategorization preferences. *Journal of Memory and Language* 48 (2): 281–303.
- Ide, Nancy and Jean Véronis
- 1998 Introduction to the special issue on word sense disambiguation: the state of the art. *Computational Linguistics* 24 (1): 1–40.
- Jorgensen, Julia
- 1990 The psychological reality of word senses. *Journal of Psycholinguistic Research* 19 (3): 167–190.
- Kaufman, Leonard and Peter J. Rousseeuw
- 1990 *Finding Groups in Data*. New York: Wiley.
- Kilgarriff, Adam
- 1993 Dictionary word sense distinctions: an enquiry into their nature. *Computers and the Humanities* 26 (1–2): 365–387.
- 1997 I don’t believe in word senses. *Computers and the Humanities* 31 (2): 91–113.
- Kishner, Jeffrey M. and Raymond W. Jr. Gibbs
- 1996 How *just* gets its meanings: Polysemy and context in psychological semantics. *Language and Speech* 39 (1): 19–36.
- Klein, Deborah E. and Gregory L. Murphy
- 2001 The representation of polysemous words. *Journal of Memory and Language* 45 (2): 259–282.

- 2002 Paper has been my ruin: Conceptual relations of polysemous senses. *Journal of Memory and Language* 47 (4): 548–570.
- Krauth, Joachim
1993 *Einführung in die Konfigurationsfrequenzanalyse: [...]*. Weinheim: Beltz, Psychologie-Verlags-Union.
- Lakoff, George
1987 *Women, Fire, and Dangerous Things*. Chicago: The University of Chicago Press.
- Langenscheidts Grosswörterbuch der Englischen und Deutschen Sprache: “Der kleine Muret-Sanders”. 1991 5th ed. Berlin/München: Langenscheidt.
- Lehrer, Adrienne
1990 Polysemy, conventionality, and the structure of the subjective lexicon. *Cognitive Linguistics* 1 (2): 207–246.
- Manning, Christopher D. and Hinrich Schütze
2000 *Foundations of Statistical Natural Language Processing*. 4th printing with corr. Cambridge, MA: The MIT Press.
- MacWhinney, Brian
2000 *The CHILDES Project: Tools for Analyzing Talk*. 3rd ed. Mahwah, NJ: Lawrence Erlbaum Associates.
- McDonald, Scott
1997 Exploring the validity of corpus-derived measures of semantic similarity. Paper presented at the 9th Annual CCS/HCRC postgraduate conference, University of Edinburgh.
- Miller, George A.
1971 Empirical methods in the study of semantics. In: Danny Steinberg and Leon A. Jakobovits (eds.), *Semantics: An Interdisciplinary Reader*, 569–585. London/New York: Cambridge University Press.
- Nordquist, Dawn
2004 Comparing elicited data and corpora: What the (mis)match reveals about the lexicon. In: Michael Achard and Suzanne Kemmer (eds.), *Language, culture, and mind*. Stanford, CA: CSLI.
- Norvig, Peter and George Lakoff
1987 Taking: A study in lexical network theory. In: Jon Aske, Natasha Beery, Laura Michaelis, and Hana Filip (eds.), *Proceedings of the Thirteenth Annual Meeting of the Berkeley Linguistics Society*, 195–206. Berkeley, CA: BLS.
- Nunberg, Geoffrey, Ivan A. Sag, and Thomas Wasow
1994 Idioms. *Language* 70 (3): 491–538.
- Oxford English Dictionary 2 on CD-ROM V. 1.15. Oxford: Oxford University Press.

Partridge, Eric

- 1961 *Origins: A Short Etymological Dictionary of Modern English*. London: Routledge and Kegan Paul.

Rice, Sally

- 1987 Towards a transitive prototype: Evidence from some atypical English passives. In: Jon Aske, Natasha Beery, Laura Michaelis, and Hana Filip (eds.), *Proceedings of the Thirteenth Annual Meeting of the Berkeley Linguistics Society*, 422–434. Berkeley, CA: BLS.

- 1996 Prepositional prototypes. In: Martin Pütz and Rene Dirven (eds.), *The Construal of Space in Language and Thought*, 135–165. Berlin/New York: Mouton de Gruyter.

Roland, Douglas and Daniel Jurafsky

- 1998 How verb subcategorization frequencies are affected by corpus choice. *ACL/COLING Proceedings*, 1122–1228.

- 2002 Verb sense and subcategorization probabilities. In: Paola Merlo and Suzanne Stevenson (eds.), *The Lexical Basis of Sentence Processing: Formal, Computational and Experimental Issues*, 303–324. Amsterdam/Philadelphia: John Benjamins.

Sandra, Dominiek

- 1998 What linguists can and can't tell you about the human mind: A reply to Croft. *Cognitive Linguistics* 9 (4): 361–378.

Sandra, Dominiek and Sally Rice

- 1995 Network analyses of prepositional meaning: mirroring whose mind – the linguist's or the language user's? *Cognitive Linguistics* 6 (1): 89–130.

Schulte im Walde, Sabine

- 2003 Experiments on the automatic induction of German semantic verb classes. Doctoral dissertation, University of Stuttgart.

Stefanowitsch, Anatol and Stefan Th. Gries

- 2003 Collostructions: Investigating the interactions of words and constructions. *International Journal of Corpus Linguistics* 8 (2): 209–243.

Theakston, Anna L., Elena V. M. Lieven, Julian M. Pine, and C. F. Rowland

- 2001 The role of performance limitations in the acquisition of verb-argument structure: An alternative account. *Journal of Child Language* 28 (1) 127–152.

- 2002 Going, going, gone: The acquisition of the verb 'go'. *Journal of Child Language* 29 (4):783–811.

Tuggy, David

- 1993 Ambiguity, polysemy, and vagueness. *Cognitive Linguistics* 4 (3): 273–290.

Tyler, Andrea and Vyvyan Evans

2001 Reconsidering prepositional polysemy networks: The case of *over*.
Language 77 (4): 724–765.

Williams, John N.

1992 Processing polysemous words in context: Evidence for interrelated
meanings. *Journal of Psycholinguistic Research* 21 (3): 193–218.