
Musings about the Impossible Electronic Dictionary

Annie Zaenen

Parc, Zaenen@parc.com

Sue Atkins and I met over discussions about the polytheoretical, multi-functional lexicon, a bit more than 15 years ago when it became obvious that the computer would have an important impact on how dictionaries were made. The impact was expected to go in two directions. On the one hand, dictionary making would be freed from several of its constraints because more data could be accessed quicker to insure better coverage, size limitations could be overcome and the organization could be more flexible because the information could be accessed in various ways. This much was applicable to human readable electronic dictionaries. On the other hand, it was hypothesized that the information contained in traditional dictionaries would be useful for natural language applications. It was also recognized that the information in traditional dictionaries would not be sufficient and that it would be an enormous task to make adequate lexical databases for NLP starting from that information.

The notion of *polytheoretical* dictionary intended to do away with at least one obstacle that was feared to be in the way of the construction of large reusable lexical databases: the differences in linguistic theories and notational conventions. Of course, we were aware of the fact that linguistic theories in general had rather little to say about lexical matters (see e.g. Zaenen and Engdahl, 1991) but the paradigm for natural language applications at that moment included sentence parsing as an obligatory step, hence the need for lexical resources to be adapted to the various assumptions that parsers might make. Through the years it has become clear that the task was even more monumental than we anticipated but also that, at the moment and most likely for some time to come, the sentence parsing paradigm is of minor interest in natural language applications because, with the availability of the internet for more and more people, the most pressing task has become information retrieval for which parsing is much too slow and unnecessary.

The interest in lexical resources for NLP has not gone away. It has led to new models of lexical organization in the symbolic processing community and to radically different proposals in the statistical community. But the evolution has shown that different applications have different needs and the idea of an extensible reusable multi-purpose lexical database has become less prominent and is certainly not closer to a practical realization now than fifteen years ago. In what follows I give a quick subjective overview of what I see as the evolution and the highlights during these fifteen years and evaluate equally subjectively the viability of the multi-functional enterprise.

The perceived need for a more explicit and precise representation of lexical knowledge has led to increased interest in lexical semantics. In the mid-eighties formal linguistic theory had close to nothing to say about the most important problems of lexical semantics. The last fifteen years have seen a great change in this situation. Both within established linguistic theories and within new frameworks the lexicon has gotten quite a bit more attention.

The efforts inspired by syntactic theories have concentrated mainly on valency alternations: different syntactic realizations of what is taken to be an 'event' with the same kind of participants, e.g. the relation between "*John gave a book to Mary*" and "*John gave Mary a book*". This includes the work on LMT in LFG (e.g. Alsina, 1993 Bresnan et al. 1989, 1990), GB (e.g. Grimshaw, 1990), Categorical grammar (e.g. Dowty, 1991), HPSG (e.g. Davis, 1995) and Construction Grammar (e.g. Goldberg, 1995). All these publications develop theories to account for valency alternations but are not concerned with large-scale descriptions. The only work that aims at descriptive coverage is Levin, 1993, an extensive but not exhaustive depository of valency classes and their semantic characterization.

Work closer to traditional semantic tradition, e.g. by Krifka (1998) and by Dowty (1991) has brought lexical semantics more in tune with formal linguistic semantics. Work here has focused especially on verbal aspect and has led to detailed explorations of the way events are talked about in natural language.

A less syntactically inspired trend has concentrated on regular sense extensions and their limits, e.g. the fact that 'chicken' can refer to an animal running around in the yard and to some piece of cooked meat on your plate. Here we find work by Nunberg (1995), Nunberg and Zaenen (1992) and Briscoe, Copestake and Lascarides (1995) who show what technical prowess is needed to distinguish between the edible chicken and the inedible pig¹.

¹ at least with knife and fork. Coyotes and the like can eat pig.

The interest for semi-productivity has been the main center of attention for lexicologists coming from the linguistic tradition. The most elaborate theoretical contribution is Pustejovsky's *Generative Lexicon*. Instead of listing sub-senses, the generative lexicon aims at making them result from the combination of the general meanings of the various words in a sentence through lexical rules that use an enriched representation for lexical items. The mechanism has been mainly used to explain logical metonymy. For instance a noun like *book* has associated with it a qualia structure specifying among other things its purpose and function and factors involved in bringing it into existence. This enriched structure is meant to explain why we can say sentences like "*He began the book.*" Although the book in itself is an object without temporal structure it has associated to it a purpose role: to be read, and an agentive role, to come into existence through writing, which have a temporal structure. These roles are exploited to derive the meanings of the example above.

The Generative Lexicon has certainly generated a lot of debate and has led to the extensive discussion of phenomena that had received little attention previously. It is, however, not clear that it has proposed a viable model for the lexicon. More detailed studies of specific lexical items suggest that the proposed qualia structure is not sufficient to account for even logical metonymy (see for instance the discussion in Jayez, 2001). But more importantly the enterprise started with a point of view that seemed to endorse a difference between meaning and use and between semantics and pragmatics and it has mainly led to a more articulate questioning of these distinctions. First, in the light of possible interpretations of examples such as "*She enjoyed the book*" as "*She enjoyed eating the book*" (if she is a goat), recent papers stress the importance of context. This has provoked many amendments to the original proposal (see e.g. Lascarides and Copestake, 1998, Copestake and Lascarides, 1997). Second, the limited productivity of logical metonymy, documented in empirical work (see Verspoor, 1997a, 1997b) has drawn renewed attention to lexical idiosyncrasy and conventionalization. The changes incorporating context and conventionalization complicate the implementations of the Generative Lexicon and rob it of much of its predictive power. Furthermore it has been shown, in an admittedly rather small-scale study (Kilgarriff, 2001), that lexical forms that are not found in traditional lexicons are not predicted by the GL approach.

The theory can take two different directions in the face of these criticisms. One is to circumscribe precisely what it intends to be a theory of in such a way that, for instance, conventionalization and lack of attested examples of the generative power of its rules are not one of its concern. At that point it will of

course not be relevant for NLP. Or it can try to extend its aims to account for context and limited productivity. Going this way will most likely change its character rather drastically but it seems to be the direction that the researchers connected to the enterprise are taking. For the moment the approach over- and undergenerates and has no clear computational way to incorporate the context factors or to limit overgeneration efficiently.

Frame Semantics proposes a more conservative approach to the lexicon and also one more rooted in traditional philological and lexicographic concerns. It does not try to predict new senses but rather to integrate and motivate them from existing meanings. It also does not try to abstract away from world knowledge. On the contrary the researchers in that paradigm assume that the lexicon is structured by the reality that underlies or frames it. Sense extensions occur from a prototypical center out. Frame semantics observes that to account for extended senses it is necessary to register what could be called different uses because the extended senses are derived from them (see Fillmore and Atkins, 2000, pp 101-102)².

Within this approach, Fillmore and Atkins and their collaborators have provided detailed analyses for word senses, for instance, *risk* (Fillmore and Atkins, 1991) and *crawl* (Fillmore and Atkins, 2000) and evaluated the performance of existing dictionaries against them. Their work has revealed important lacunae in the coverage of existing dictionaries. The lacunae concern missing word senses but also basic syntactic information that the type of dictionary that Fillmore and Atkins (2000) concentrated on, the monolingual dictionary for non-native speakers, promises to provide. Bilingual dictionaries are shown to be not more complete and tend to have even less syntactic information.

The Berkeley FrameNet project attempts to systematize and automate the creation of lexical entries based on conceptual frame structures. It could have an important impact on dictionary construction per se but it is unclear that it will because the effort involved seems to exceed what any dictionary publisher would be willing to invest. For NLP applications, this work suffers from the opposite problem of the one that afflicts the Generative Lexicon: it does not generate novel uses. It lays the basis for such work but needs an explicit proposal to handle metaphor and metonymy computationally.

² The lexical work is associated with a more general theory, Construction Grammar that, contrary to most work on sentence semantics (most explicitly the Montague Grammar tradition), attributes an important role to syntactic constructions as bearers of meaning (see e.g. Goldberg, 1995, 1997).

The last effort that has shaped much of the reflection about the lexicon in the English-speaking world is *WordNet*. It started from dissatisfaction with the traditional dictionary as a means to help people, and especially children, understand words. It was conceived by George Miller and developed by Fellbaum and other of his associates (see Fellbaum, 1998, for a collections of articles on the project). The initial inspiration was psychological but the idea goes back to the old thesaurus concept of which WordNet is a sophisticated variant. Over the dictionary it has the advantage to represent the hierarchical relations between lexical items explicitly but it has a rather uneasy status between reorganized dictionary and ontology. And although originally Miller's dissatisfaction with traditional dictionaries was that they did not provide context to allow children to understand what the definitions meant, WordNet doesn't represent context anymore than a traditional dictionary.

Being initially based on English only, the model has been extended to other languages but these extensions are too small to play an important role in NLP. The English WordNet, however, has become the lexical initiative most exploited in NLP and has taken over the role previously played by machine readable dictionaries: everybody tries to use it because one is convinced that it must be better than starting from scratch. One can doubt this seeing the efforts spent in adapting WordNet to the need of various projects while still having to cope with its unavoidable gaps and shortcomings but its hierarchical structure makes it an important tool to help overcome the sparseness of data problem of statistically based lexical acquisition: it allows one to create equivalence classes e.g. on the basis of hypernyms. This is much more laborious when one has to start from a dictionary (see for instance, Leacock et al, 1998).

Let me also mention a lexical theory that has, most likely for socio-cultural reasons, less impact on NLP in the US but has some influence in Europe and Canada: the *Meaning-Text* model developed by Mel'čuk and his collaborators. It has among other things developed a notion of lexical function that makes it easy to describe certain types of collocations. For instance, intensity is often marked in idiomatic ways: *rain is heavy* but *cold is bitter*. Lexical functions permit to collect these collocations a limited number of types that exist across languages and to associate if necessary cross-linguistic expressions with them. In French the rain is strong (*une forte pluie*) and it is duck cold (*froid de canard*).

The model has been used in the construction of a partial French dictionary: *Dictionnaire explicatif et combinatoire du français contemporain*. (Mel'čuk et al. 1984, 1988, 1992, 1999). The current vogue of dependency grammars of

which the Meaning-Text model is an instance will most likely lead to more interest in the future.

While lexicology was addressing the issues discussed above, more application-oriented approaches have evolved in two different directions (which doesn't exclude combinations): statistical modeling and knowledge representation. Neither of those takes a traditional lexical database as central and neither distinguishes sharply between meaning and use, although they exploit traditional dictionaries among other sources of information.³

Reluctantly or not, many researchers confronted with the need of constructing large-scale lexicons seem to be coming to the view that concepts need to be represented as first class citizens in the lexical enterprise. We see a renewed interest in *ontologies* that make a distinction between language particular words and language independent concepts. This point of view has been advocated in the machine translation community for a long time for the obvious reason that different languages lexicalize concepts in different ways. As soon as one starts to think about more than two languages, the interlingua approach starts to look promising. When reasoning is important the need to go beyond words to concepts is also evident even in a monolingual context.

Nirenburg (see e.g. Nirenburg and Raskin, 2001) and his coworkers for instance have been laboring at a knowledge-based machine translation project with a handcrafted ontology for more than ten years. Translation is a practical context in which the generative character of the lexicon could be exploited but one sees that translation specialists stress the great diversity of types of generativity and the limits of each specific type, which make generativity as such unfortunately a notion of limited practical importance.

Ontology will most likely gain more attention with the rise of the semantic web. This is still an embryonic enterprise and only very abstract specifications are discussed but it will be interesting to see whether it can have influence beyond the domains in which a disambiguated controlled language can be imposed.

The most important development of the last ten years is, however, the rise of *statistical models*. In statistical approaches the use of a word IS its meaning. This is the clearest in extreme versions of the approach, for instance represented by the work of Hinrich Schütze (1997, 2000) who sees the sense of a word as 'a group of contextually similar occurrences of a word'. Schütze does away with all recourse to dictionaries to treat ambiguous words. Meaning is completely

³ For an overview of the use of machine readable dictionaries, see Wilks et al. (1996)

reduced to use via co-occurrence calculations, in practice of second order, in theory of any order. He sidesteps accordingly the vexing problem of the number of senses that need to be distinguished: it will be determined by the number of clusters allowed. Vagueness and ambiguity, homonymy and polysemy are not distinguished here except as poles of a continuum. The idea that there is no need for sense labels or even for discrete senses and that what corresponds to word senses can be represented by co-occurrences is evidently gaining ground. Recent elaborations of **MindNet** embody the idea in the form of a spreading activation network. One might wonder though how much refinement the methods will need for idiomatic expressions not to show up as noise (e.g. *chercher midi à quatorze heures*, to look in the wrong place) or to be completely ignored (e.g. *chercher la petite bête*, to be overcritical, literally: to look for the little animal).

Can one plausibly use such an approach for other NLP tasks than IR? Given enough similar texts in different languages, it is conceivable that translation could exploit such techniques, and approaches such as IBM statistical translation model are very close in spirit, if not in implementation detail. It is, however, more difficult to see how these representations can be used in reasoning, except of course if propositional reasoning is also replaced by some kind of associative reasoning.

From a practical point of view, it is not clear that the statistical specification of the sense disambiguation task makes it more feasible. Statistical data are very corpus sensitive and no corpus seems ever to be big enough. Less revolutionary statistical approaches use similarity measures for sense disambiguation, where the various senses are given by an existing machine-readable dictionary. These methods have evolved from supervised learning (with a hand tagged training corpus) to unsupervised learning (Yarowsky, 1995) or by doing the tagging automatically (Karov and Edelman, 1998). While the results of these experiments are encouraging, it has to be kept in mind that, like Schütze's, they typically try to disambiguate words that are more often homonyms than cases of real polysemy and work best given rather rough clusters. Given the observations of Atkins about the lacunae in existing dictionaries it is doubtful that very fine distinctions can be made when a dictionary is used as an essential source.⁴

⁴ It is of course not the case that humans always understand perfectly what they hear or read. For instance when I first read in a contribution to a distribution list "*I have been lurking in the background for a while*", I did not read this as a special use of 'lurk'. It was only after a couple of mentions that it dawned on me that what I took for a creative metaphor was in fact a specialized use of the word.

Another interesting and completely unforeseen development has been the role played by the web. Whereas in earlier days of NLP, the dictionary was used as a small but fairly well structured corpus, the availability of textual data and the statistical methods of the type alluded to above make it possible to see any corpus as a lexical database in a fairly direct way. This is especially true of the web : combined with more or less sophisticated statistical tools, it can take over several of the functions of the traditional dictionary⁵. If one doubts about the spelling of a word, one can type in the various versions and, relying on the possibly dubious assumption that a word is more often spelled right than wrong, see which one wins. When one doesn't know the meaning of a word, the Web gives, even without special tools, a quick and dirty concordance that in many cases is sufficient for one's needs. In a bilingual context, the web provides easy assistance with terminology and collocations: if you doubt whether the translation of '*propositional logic*' in French is '*logique propositionnelle*' or '*logique des propositions*', the web is most likely a better place to look than a general dictionary (but again one should take the ratio of mentions into account, both expressions are found). Less anecdotally, the Web has been exploited to make bilingual lexicons with text alignment methods. For instance, the EC project Twenty-One exploited web resources to make its dictionaries, based on alignment statistics (see the Twenty-One home web page for the results).

While this kind of exploitation of the web is tempting, one has to keep in mind that the web is very skewed and extremely time sensitive. About a year ago I searched for the expression *jeune pousse* to see whether it had been accepted as a translation for *start-up*. While doing this I found another sense (or use) of *jeune pousse* that was not in my dictionary: a budding entertainer or sportsperson, as illustrated in the following context "*Alizée, cette jeune pousse de la chanson française dont le tube Lolita est encore dans toutes les mémoires*"⁶. At that moment, the sense jumped out of the web pages. When I did a new search a couple of weeks ago with the same search engine, I could only find it because I knew what I was looking for. It had been completely overwhelmed by start-ups and the traditional horticultural sense.

What the Web also doesn't have is the aura of authority and the power to decide what is right and what is wrong that people ascribe to the traditional dictionary. Even if one thinks that the authoritative value of the traditional

⁵ A version of some of the remarks that follow was presented by Gregory Grefenstette in oral presentations at XRCE. See also his article pp. 199-215 of this book.

⁶ Alizée, the young promise of the French song whose hit Lolita is still in everybody's memory...

dictionary is overestimated, it is clear that, in some contexts, more is needed than (co)occurrence statistics.

An illustration can be found in so-called forensic linguistics. A couple of years ago, my colleague Geoff Nunberg was asked what the meaning of the word *broadcast* was between 1940 and 1960. The important point of the discussion was whether the word might have implied point-to-point transmission of information. Nunberg (1998) argued that it didn't, partly on the basis of evidence about the way the word acquired its media sense in the twenties and the way it behaved more recently. At first it was used to distinguish radio transmission from point-to-point transmission by telegraph. The word was apt to embody this contrast because its original meaning as sowing by scattering seed over a whole surface. Radio waves are undirected and can be picked up by anybody who has the right instrument in contrast with telegraphy that is directed to a particular destination. Given new technology such as cable, the question is now whether *broadcasting* can also be used for these ways of transmitting non-written information to mass audiences. The cable model is opposed to the original radio model in just the two crucial ways that telegraphy was opposed to it; it is point-to-point and not free. Had the word, so to speak, forgotten its origins and could it now be generalized to any form of transmission of news and entertainment? A careful examination of corpus data from recent newspaper citations shows that this is not the case and, that, although there are cases in which *broadcast* is used to cover *cablecast*, in general *broadcasting* and *cable transmission* are kept distinct. There is a *Journal of Broadcasting and Cable*, there are awards for 'broadcasting and cable excellence', etc. So *broadcast* doesn't seem to cover *cablecast*, except in the way that the use of *mouse* occasionally covers all devices with a similar function. The fact that the sense extension, which could have occurred, didn't occur is the strongest evidence that the word *broadcast* had in a previous period also the more restricted meaning. The same argument can for instance not be made about the relation between old-fashioned and digital watches: as soon as the digital watches appeared they were obviously *watches*. This kind of subtle lexicographic reasoning, like the subtle representations of related word senses in frame semantics, cannot be reduced to corpora statistics alone.

It is clear then that the multi-purpose dictionary is not playing the expected role in NLP and that in fact its importance has diminished. Our other assumption was that computational tools could lead to a substantial improvement in the quality of traditional dictionaries. This hasn't happened either.

Computational tools have turned out to be very valuable for dictionary makers. The time of the shoeboxes is beyond us and dictionaries are maintained and formatted in electronic version and made on the basis of electronically available corpora. In the Anglo-Saxon world large corpora have been constructed that can be exploited for lexicon construction, most notable the British National Corpus.⁷ Concordance tools allow the data to be sorted quickly. Statistical techniques that go beyond occurrence ranking are also used as aids to lexicon construction. For instance statistical alignment of parallel text in two languages is an important aid in the development of terminological and other lexical databases.

But while the work by e.g. Atkins has shown how machine-readable resources can be exploited to discover and to overcome the limits of traditional dictionaries, it nor any other of the lexical investigations mentioned above seem to have had much impact on actual dictionary content. Maybe this is just due to the fact that institutions like dictionaries change slowly but it might also be the case that they simply do not feel the need to improve in the ways that recent research would make possible. Dictionaries are cultural artifacts and as such they respond with the available technology to the social needs of a community at a given time. The typical needs that a traditional monolingual dictionary serves are checking the spelling or the meaning of an occasional word and in most cases, to inform about what the community considers to be acceptable words and meanings and which are not. It is certainly not used as the representation of all the knowledge about words of a language community and it is not meant to be consulted word after word to decipher the meaning of texts. The idea of an all-compassing dictionary is a 19th century ambition superimposed on a basically utilitarian or regulatory enterprise started some centuries earlier. The 19th century impulses for exhaustive classification have in fact given most people a rather skewed idea of what a language is and have given the dictionary the status of a kind of Bible that it never could completely live up to. Changing the format would jeopardize this status and is thus an enterprise that will be undertaken very cautiously. Given the way the dictionary is used, it is only if the current format is seen as an impediment to its social status that it will be given up. Other adaptations will be minor and gradual.

There is one type of monolingual dictionary -- which exists for the moment only for English -- that should imperatively take into account and extend the

⁷ The situation is less good in other language areas, for instance in French there is no national large scale initiative to construct a representative corpus that goes beyond literature and access to what exists is difficult if one does not belong to the right type of institution.

findings of Atkins and her coauthors: the learner's dictionary. These dictionaries are consulted by non-native speakers to get precise information about the use of words, their syntactic patterns and their collocations. Better information and a more adequate structure of the information should enhance their utility in important ways. It is, however, unlikely that they constitute a market that can sustain the investment needed for these improvements.

Looking back at the last ten fifteen years then, it is obvious that the field has developed in rather different ways from those anticipated. The machine readable improved dictionary does not have a central place in word-based NLP. There are several reasons for that. One is that with the spread of the Internet, information retrieval has become such an all-overwhelming concern that its needs are providing the new paradigms for NLP. Although there seems to be a shift to tasks like question answering, it is likely that the requirements for speedy rather than broad sense disambiguation will remain a dominant concern for some time to come and that statistical methods, possibly using some more structured input than pure bags of words, will remain the most used.

There remains, however, an unmet need for more sophisticated disambiguation techniques. Optimistic reports relegating lexical ambiguity to domains such as creative writing or free wheeling speech, have proven futile. It is as much of a problem in scientific or technical writing as it is in more mundane genres. For instance, the bioinformatics sector would like to associate the knowledge contained in scientific papers with that available in dedicated databases for genes or proteins. This effort is mightily hampered by the ambiguous ways in which scientists refer to genes and any other objects in their domain of discourse. Bioinformatics is not exceptional in that respect and even in the domain of technical documentation where there are some explicit attempts to control ambiguity there remain important problems once one tries to go from text to reasoning. Working on copier maintenance manuals I spent last month an enjoyable couple of hours discussing the ambiguities (or at least the complexities) of *excessively*.

If we look at the original tasks envisioned for the multi-functional lexicon, which were tasks that needed fine-grained distinctions for applications such as good quality translation, it is also clear that we vastly underestimated the scope of the task. Especially important factors that were not taken into account are the speed with which the lexicon changes and the pervasiveness of metaphor and metonymy. These cannot be ignored but they can also not be predicted. Without the Hundred Years War, English gardens would most likely not *swarm with bees* (see Dowty 2000) and without a song and a commercial we would have had

many less *bad hair days* over the last years. More importantly even with the Hundred Years War and the song and the commercial we might not have had *gardens swarming with bees* or *bad hair days*. In a 1992 paper Geoff Nunberg and I reiterated an old observation that we need two kinds of rules, one kind that concerns itself ‘only with a repertory of transfer functions provided by pragmatics or highly general semantic principles’ and another kind about ‘licenses and conventions of use of the speech community’, which are not predictable on linguistic grounds.

Work grounded in the linguistic tradition of the second half of last century has ignored the cultural and encyclopedic side of things and the integration of lexical semantics with formal semantics will only lead to a further neglect of these aspects. If such approaches are not supplemented with some that do take culture and society into account, we will end up with interesting insights in many general rules but with little insight in why these rules account for so few instances and with no explanation for a lot of things that do occur. For example, an approach based on linguistic tradition will have nothing to say about why we understand the following 2 passages which I just lift out of *Libération* of the last days. “*Fierté voire chauvinisme, cultures incompatibles, fâcheuses réalités après la chaleur du banquet, rien de particulier à l'exemple de DaimlerChrysler, rien de nouveau dans le kamasoutra des entreprises.*” (*Libération* 6/4/01) “*Le chanvre roule carrosse Il pourrait détrôner la fibre de verre dans la fabrication des autos.*” (*Libération* 6/2/01)⁸ A linguistically anchored theory does not need to have anything to say about such sentences but a NLP system and a foreign language learner cannot ignore all such occurrences. Hence they will need other sources of knowledge than those provided by linguistic resources alone⁹. Frame semantics is more responsive than most approaches to these concerns. If worked out further it could be a great source of insights in how words acquire their uses and senses and it could be of great help for foreign language learning. Would it, however, have much impact on NLP? It is clear that the lexical insights that it can deliver do not translate directly into computational models. The type of productivity that we see at work

⁸ “Pride even chauvinism, incompatible cultures, unpleasant realities after the warmth of the banquet, nothing special in the example of DaimlerChrysler, nothing new in the Kamasutra of enterprises”

“Hemp is doing well. It could replace glass fiber in the construction of cars”. The point here is that the expression chosen to say ‘doing well’ is *rouler carrosse*, meaning literally “to be transported in a coach.”

⁹ Note that even for simpler looking examples we have similar problems: how does one predict that *paint program* and *draw program* will be distinct exactly in the way they are?

in natural language use can most likely not be modeled directly with the methods that we dispose of now, although some attempts have been made based on Cognitive Semantics (see e.g. Narayanan 1997)

Contrary to the aim of a unified lexical database it seems that we will have several different enterprises: traditional dictionaries, rather small lexical databases that are structured according to the new insights of the various lexical theories that have developed in the past ten years, bigger lexicons for various applications, most of them not human-readable and based on statistical methods. Given the computational tools that most researchers now have at their disposal, even the more theoretical studies, which do not aim at the development of large lexicons, however, will be accountable for the facts and will test out their hypotheses via implementations. That big step forward in linguistic theory formation is already illustrated in projects such as the Generative Lexicon and FrameNet.

Bibliography

- Alsina, Alex, 1993, *Predicate Composition: A theory of Syntactic Function Alternations*, PhD thesis, Stanford University, Stanford CA
- Bouillon, Pierrette and Federica Busa, eds., 2001, *The Language of Word Meaning*, Cambridge University Press, Cambridge UK.
- Bresnan, Joan and Jonni Kanerva, 1989, Locative Inversion in Chichewa, *Linguistic Inquiry* 20
- Bresnan, Joan and Liobi Moshi, 1990, Object Asymmetries in Comparative Bantu Syntax, *Linguistic Inquiry* 21
- Briscoe, Ted, Anne Copestake and Alex Lascarides, 1995, Blocking, in P. St. Dizier and E. Viegas (eds), *Computational Lexical Semantics*, Cambridge University Press, pp.273-302
- Briscoe, Ted and Anne Copestake, 1996, Controlling the application of lexical rules, in *Proceedings of the ACL SIGLEX Workshop on Breadth and Depth of Semantic Lexicons*, Santa Cruz, CA.
- Copestake, Ann and Alex Lascarides, 1997, Integrating Symbolic and Statistical Representations: The Lexicon Pragmatics Interface, in *Proceedings of the 35th Annual Meeting of the ACL and the 8th Conference of the EACL*, pp.136-143
- Davis, Anthony, 1995, *Lexical Semantics and Linking in the Hierarchical Lexicon*, PhD. Thesis, Stanford University, Stanford CA.
- Dolan, William, Lucy Vanderwende and Stephen Richardson, 2000, Polysemy in a broad coverage Natural Language System, in *Ravin and Leacock*.
- Dowty, David, 1991, Thematic Proto-Roles and Argument Selection, *Language*, 67 (3)
- Dowty, David, 2000, The Garden is Swarming with Bees, in *Ravin and Leacock*.
- Fellbaum, Christiane (ed), 1998, *WordNet: An Electronic Lexical Database*, MIT Press, Cambridge MA

- Fillmore, Charles and B.T.S. Atkins, 1991, *Toward a Frame-Based Lexicon: The Semantics of Risk*, Proceedings of the 29th ACL, Berkeley CA
- Fillmore, Charles and B.T.S. Atkins, 2000, *Describing Polysemy: The Case of "Crawl" in Ravin and Leacock*
- Goldberg, Adele, 1995, *A Construction Grammar Approach to Argument Structure*, University of Chicago Press, Chicago
- Goldberg, Adele E., 1997, *Construction Grammar*, in E.K. Brown and J.E. Miller (eds.), *Concise Encyclopedia of Syntactic Theories*. New York: Elsevier Science Limited.
- Grefenstette, 2001, oral presentation at XRCE, Grenoble
- Grimshaw, Jane, 1990, *Argument Structure*, MIT Press, Cambridge MA.
- Jayez, Jacques, 2001, *Underspecification, Context Selection and Generativity*, in Bouillon and Busa
- Karov, Yael and Shimon Edelman, 1998, *Similarity-based Word Sense Disambiguation*, *Computational Linguistics*, 24,1.
- Kilgarriff, Adam, 2001, *Generative lexicon meets corpus data: the case of non-standard word uses*, in Bouillon and Busa
- Krifka, Manfred, 1998, *The Origins of Telicity*, in S. Rothstein (ed), *Events and Grammar*, Kluwer Academic Publishers, Dordrecht NL
- Lascarides, Alex and Ann Copestake, 1998, *Pragmatics and Word Meaning*, in *Journal of Linguistics*, 34.2, Cambridge University Press.
- Leacock, Claudia, Martin Chodorow and George Miller, 1998, *Using Corpus Statistics and WordNet Relations for Sense Identification*, *Computational Linguistics*, 24, 1.
- Levin, Beth, 1993, *Towards a Lexical Organization of English Verbs*, University of Chicago Press, Chicago, IL
- Mel'čuk, I. A., N. Arbatchewsky-Jumarie, L. Elnitsky, L. Iordanskaja and A. Lessard. 1984. *Dictionnaire explicatif et combinatoire du français contemporain: Recherches lexico-sémantiques I*, Montréal, Presses de l'Université de Montréal, 172 pages.
- Mel'čuk, I. A., N. Arbatchewsky-Jumarie, L. Dagenais, L. Elnitsky, L. Iordanskaja, M.-N. Lefebvre and S. Mantha. 1988. *Dictionnaire explicatif et combinatoire du français contemporain: Recherches lexico-sémantiques II*, Montréal, Presses de l'Université de Montréal, 332 pages.
- Mel'čuk, I. A., N. Arbatchewsky-Jumarie, L. Iordanskaja and S. Mantha. 1992. *Dictionnaire explicatif et combinatoire du français contemporain: Recherches lexico-sémantiques III*, Montréal, Presses de l'Université de Montréal, 323 pages.
- Mel'čuk, I. A., N. Arbatchewsky-Jumarie, L. Iordanskaja, S. Mantha and A. Polguère. 1999. *Dictionnaire explicatif et combinatoire du français contemporain: Recherches lexico-sémantiques IV*, Montréal, Presses de l'Université de Montréal, 347 pages.
- Narayanan, Srin, 1999 *Moving Right Along: A Computational Model of Metaphoric Reasoning about Events* Proceedings of the National Conference on Artificial Intelligence, AAAI Press.
- Nirenburg, Sergei and Vicor Raskin, 2001, *Choices for Lexical Semantics*, *Computational Intelligence*, 17
- Nunberg, Geoff, 1995, *Transfers of Meaning*, *Journal of Semantics*, 12

- Nunberg, Geoff, 1998, *The Meaning of the Word Broadcast, 1940-1960*, unpublished paper PARC, Palo Alto CA
- Nunberg, Geoff and Annie Zaenen, 1992, *Systematic Polysemy in Lexicology and Lexicography*, in H. Tommola et al. (eds) *Proceedings of Euralex II*, University of Tampere, Finland
- Pustejovsky, James, 1995, *The Generative Lexicon*, MIT Press, Cambridge MA
- Ravin, Yael and Claudia Leacock (eds), 2000, *Polysemy*, Oxford University Press, Oxford UK
- Schütze Hinrich, 1997, *Ambiguity Resolution in Language Learning*, CSLI Publications, Stanford CA
- Schütze Hinrich, 2000, *Disambiguation and Connectionism*, in Ravin and Leacock *Twenty-One Lexicons*: <http://wwwhome.cs.utwente.nl/~irgroup/align/>
- Verspoor Cornelia M., 1997a, *Contextually-Dependent Lexical Semantics*, PhD thesis, The University of Edinburgh, Edinburgh
- Verspoor, Cornelia M. 1997b, *Conventionality-Governed Logical Metonymy*, in H.Bunt et al. (eds), *2nd International Worskop on Computational Semantics*, Tilburg, NL
- Wilks, Yorick, B. Slator, and Louise Guthrie, 1996, *Electric Words*, MIT Press, Cambridge, MA.
- Yarowsky, David, 1995, *Unsupervised word sense disambiguation rivaling supervised methods*, *Proceedings of the 33rd ACL meeting*, Cambridge MA
- Zaenen, Annie and Elisabet Engdahl, 1991, *Descriptive and Theoretical Syntax in the Lexicon*, in B.T.S. Atkins and A.Zampolli (eds), *Computational Approaches to the Lexicon: Automating the Lexicon II*, Oxford University Press, Oxford, UK.