Computational Lexical Semantics in Review

Ted Briscoe
ejb@cl.cam.ac.uk
Computer Laboratory
University of Cambridge
Pembroke Street
Cambridge CB2 3QG, UK

1 Introduction

In this short overview, whose title I have shamelessly adapted from Beth Levin's (1985 (1995)) influential review of linguistic work on (verbal) lexical semantics, I try to rereview this and related work in the light of the last decade's research and also to identify developments relevant to computational lexical semantics. In the final section, I summarise how the Acquilex (Esprit Basic Research) project has attempted to extend and exploit this work in the development of a prototype computational lexicon incorporating lexical semantic information which integrates with computational approaches to proof/model-theoretic compositional semantics (of the type explored in the FraCaS (CEC LRE) project).

The study of lexical semantics and of the lexicon has seen a resurgence of activity during the last decade. Hudson (1986 (1995)) convincingly argues that linguistic theory provided little justification for the then prevalent mainstream assumptions that:

1. the lexicon is a distinct component of the grammar,

2. the lexicon is a discrete list of lexical entries, and

3. the lexicon contains only intralinguistic information.

Assumption 2) has been largely abandoned and replaced by a conception of the lexicon as an inheritance network incorporating at least is-a links (e.g. Pollard & Sag, 1987) and possibly default inheritance (e.g. Evans & Gazdar, 1989a,b). One important implication of this development for lexical semantics is the possibility of treating the polysemy/homonymy distinction more satisfactorily by treating polysemy in terms of such links rather than ignoring it and enumerating all word senses as list of distinct lexical entries (e.g. Pustejovsky & Boguraev, 1993). Assumption 1) is the subject of much more focussed debate as a result of the work on network-based lexical knowledge representation languages (LKRLs): in particular, Gazdar & Evans propose that whilst syntax and compositional semantics can be expressed within a monotonic unification-based formalism, lexical processes and relations require a more expressive non-monotonic formalism such as their path equation LKRL, DATR; whilst, within HPSG, it has been proposed that the lexicon is represented as a set of typed feature structures related by subsumption, effectively reducing the lexicon to the syntactic/semantic formalism. However, the status of lexical rules within the HPSG framework has been contentious and problematic (e.g. Riehemann, 1993; Calcagno, 1995). Assumption 3) has received less attention, but nevertheless can be considered in a more focussed manner in the light of the advent of LKRLs. Both DATR and the HPSG lexicon are couched in formalisms which allow only a proper subset of the operations defined in general purpose KRLs which have been proposed within the AI literature. It remains to be seen whether lexical relations and processes can be adequately expressed within these more restricted formalisms.

These brief remarks, I hope make it clear that lexical (semantic) work has progressed somewhat beyond the situation which Hudson criticises and which he exemplifies with the GPSG (Gazdar *et al*, 1985) lexical entry for *weep* repeated below:

```
<weep,
[[-N,[+V],[BAR 0],[SUBCAT 1]],
{wept},
WEEP'>
```

in which one sense and syntactic realisation is represented by a minimal description of orthography/phonology, syntactic (sub)category, irregular morphological variant form, and unanalysed semantic primitive, as part of a large list of unrelated other entries. In what follows I attempt to flesh out these developments in more detail, I will review lexical (semantic) research under the four subheadings of *argument structure*, *event structure*, *qualia structure* and *inheritance structure*. These are borrowed from Pustejovsky (1991, 1995), whose work has done much to rekindle and refocus work on lexical semantics. I don't attempt to be at all comprehensive either in coverage of active topics of research or in citing relevant work, rather I attempt to identify the important ideas and give main references from which further work can be found.

1.1 Argument Structure

Work on argument structure, defined as a specification of the number, syntactic and semantic type, and syntactic realisation of a predicate's arguments, has been the focus of most generative research on lexical semantics and particularly on the syntactico-semantic interface. Levin (1985 (1995)) exclusively reviews such work which primarily addresses the related problems of linking (capturing generalisations concerning the syntactic realisation of different arguments) and of diathesis alternations (the same argument may be realised differently or omitted). For example, in *John broke the vase break* appears with two arguments, an agent and a theme / affected object. In general, agents if realised are subjects and themes are objects, as here. However, as *break* can undergo the causative/inchoative alternation, the theme can be realised as subject *The vase broke*. These facts are seen as lying at the syntactico-semantic interface since they appear to be motivated by aspects of the lexical semantics of the predicate. For instance, *break* and *eat* both undergo (in)transitivity alternations (*John ate (chips)*) but only *break* undergoes causative/inchoative because it is a unaccusative change-of-state verb. Frequently, alternations modify the semantic class or entailments of a predicate, for example *slide* is extended from a change-of-position to a change-of-possession verb when it undergoes the dative alternation – *John slid the computer to the door / John slid *the door/Bill the computer*. Such phenomena and treatments of it have been exhaustively surveyed in Levin (1993).

Attempts to express rules of linking and alternation have either been couched in terms of case/semantic/thematic roles (e.g. Fillmore, 1968) or in terms of predicate decomposition (e.g. Carter, 1976). One problem with both approaches from the perspective of computational lexical semantics is that usually neither approach is related to a proof-theoretic framework. One exception is the work of Dowty (1989, 1991) which treats thematic roles as bundles of entailments prototypically associated with predicates when the associated argument is realised. This captures both the fact that alternations typically affect meaning and that the entailments associated with a thematic role, such as agent, are not uniform across all predicates (e.g. volitionality *The drunk collided with the lamppost / The drunk wanted a scotch*). It also supports a prototypic semantic account of linking in which, for example, the argument associated with the greatest number of prototypically agentive entailments is realised as subject. Dowty modifies Davidson's event-based semantic framework to include proto-thematic roles as predicates relating event variables to participants denoted by arguments. The representation of *John broke the vase* would be roughly: $\exists(e,x) \wedge\ vase(x) \wedge\ break(e) \wedge\ p-agt(ej) \wedge p-pat(ex)$.

1.2 Event Structure

Event structure refers to the (Vendlerian) classification of predicates into states, processes and transitions (achievements or accomplishments). Dowty (1979) provides a decompositional account of event structure in terms of the primitives BE, BECOME, CAUSE, and DO. He provides a model-theoretic interpretation for three of these primitives. His and further work has explored the syntactico-semantic consequences of these distinctions relating to aspect, postmodification, and diathesis alternation and refinements and extensions have been proposed (eg. Moens & Steedman, 1988; Pustejovsky, 1995). For example, as (1a,b) illustrates, durative PP postmodifiers can modify processes, but when used with transitions must be interpreted as modifying the resulting state. This can be captured in an event-based framework which posits subevent(ualities) of running and being at home broadly of the type proposed by Dowty, as sketched in (1c).

(1) a John ran for an hour.
b John ran home for an hour (at lunchtime).
c $\exists(ee') \wedge\ run(ej) \wedge\ become(e') \wedge\ at-home(e'j) \wedge$
$for(e'1h)$
d John ran (to the store)

Similarly, (1d) illustrates that a process can be coerced into a transition by the addition of a PP argument denoting a bounded path. Since the addition of a PP is often treated as a diathesis alternation, this shows that there can be a close connection between the phenomena of argument structure and event structure. Pustejovsky & Busa (1995) push this further by arguing that subevent headedness in transitions, that is whether the process or resultant state is the head, determines which subevent argument structure is realised syntactically. Many transition predicates are inherently marked for left/right headedness but when this is unspecified the resultant ambiguity over syntactic realisation leads to the causative/inchoative alternation. Pursuing the implications of contextual effects on aspectual interpretation has lead some to argue that there are no inherent lexically-encoded distinctions and that the appropriate interpretation is entirely the result of the interaction of tense/aspect marking, choice of construction, modifiers and extralinguistic context (e.g. Zaenen, 1993). Nevertheless, some type of subevent decomposition still seems needed in order to account for scopal effects with modifiers, and so forth.

1.3 Qualia Structure

Pustejovsky (1991, 1995) attempts to integrate most work on lexical semantics into the framework of generative lexicon theory. However, the most innovative component of this theory, extending work by Moravsik (1975), is his proposal concerning qualia structure as an alternative to purely relational or decompositional descriptions of word senses. There are four qualia, *constitutive*, the relation between an object and its constituent parts, *formal*, that which distinguishes it within a larger domain, *telic*, its purpose or function, and *agentive*, factors involved in its origin. For instance, an entry for *book* might specify that it is constituted of text, its formal properties are that it consists of bound pages, a CD-ROM or whatever, its purpose is to be read and it was created by writing.

Qualia structure serves several roles in the theory. However, perhaps the main advantage of this approach to the specification of lexical meaning is that it provides a natural mechanism for co-specification of related polysemic senses. Thus, in examples like *he picked up and began to browse one of the books on the table pick up* selects the formal physical dimension and *browse* the constitutive (or possibly telic) dimension. The fact that the two predicates are readily conjoined suggests that an approach which enumerates distinct polysemic senses of this type as separate lexical entries will be inadequate.

Qualia structure in nominals is also utilised in accounts of adjectival modification (*fast motorway, fast waltz*) and logical metonymy (*John began a book/cigarette*). Verbal qualia structure is utilised as a means of representing subevents in transitions; for example, the formal role of *break* is the resultant state of being broken, whilst the process of breaking which brings this about is treated as the agentive role. It is not clear to me that this extended use of qualia structure can have an identical semantics to the nominal cases, nor that it makes much sense to talk of the purpose or physical properties of eventualities in general. A more general question concerning the representation of word sense in terms of qualia is determining the semantic status of the various dimensions in a given utterance. Intuitively, it seems that one or more dimensions are asserted and the remainder presupposed in any given usage. However, this question has not been addresssed thoroughly to date.

1.4 Inheritance Structure

Inheritance structure is defined in terms of is-a or subtyping relations between elements of lexical entries. For example, the lexical entry for *novel* will inherit elements from *book* but will specialise the constitutive dimension from text to narrative. Inheritance allows for generalisation and economy in the statement of the lexicon as well as capturing relations between entries. The precise form of inheritance allowed is intimately connected to the LKRL adopted: Pustejovsky & Boguraev (1993) assume a fully general language such as KL-ONE in which multiple conflicting default inheritance is possible; Pustejovsky (1995) utilises subsumption in typed feature structures, thus limiting himself to multiple orthoganal (non-conflicting) non-default inheritance, as in HPSG (Pollard & Sag, 1987; Carpenter, 1992); Copestake & Briscoe (1995) motivate the use of multiple default inheritance in typed default feature structures for lexical semantic representation; Gazdar & Evans (1989a,b) utilise multiple default inheritance of path equations yielding a system with similar expressive power, though reentrancy between paths in a feature structure cannot be directly manipulated.

In general, systems which make use of simpler inheritance schemes need another formal operation, usually dubbed a lexical rule, which generates derived lexical entries from basic ones. Lexical rules have been used to generate inflectional, derivational, diathesis alternated, and sense extended variants of basic entries. However, in most cases alternative analyses have been proposed which utilise underspecified representations and/or inheritance to generate the required variants. Pustejovsky & Boguraev argue that a system with the expressive power of KL-ONE allows for guided inference to lexically related concepts; for example, from the agentive dimension of *prisoner* 'imprison' it is possible to follow links to 'escape' and so forth generating a space of related and potentially relevant concepts. This step – from constrained inference in terms of is-a/default links between linguistically significant syntactic and semantic information, in the sense that it directly affects syntagmatic processes, to open-ended inference in terms of concepts – represents the bound-

$$\begin{bmatrix} \mathbf{t}^1 & \\ \mathrm{F} : & \begin{bmatrix} \mathrm{G} : \mathbf{a}/\boxed{2} \\ /\mathrm{H} : \boxed{2} \end{bmatrix} \\ \mathrm{H} : & \mathbf{b}/\mathbf{c} \end{bmatrix}$$

Figure 1: An Abbreviated Typed Default Feature Structure

ary between linguistic and real-world knowledge and inference. To the extent, that clearly linguistic lexical processes can be dealt with using a constrained framework such a distinction acquires more justification.

2 Lexical Semantics in Acquilex

The Acquilex projects have in part been concerned with the development of a computationally-tractable thoery of the lexicon focussing on lexical semantics. In what follows, I summarise relevant publications and give pointers for further reading. The papers I reference are mostly available electronically at `http://www.cl.cam.ac.uk/Research/NL/acquilex/acqhome.html`.

2.1 Typed Default Feature Structures

The LKRL we utilise is an extension of the HPSG scheme to Typed Default Feature Structures (TDFSs). TDFSs are 'double' typed feature structures in which the second defeasible feature structure must be subsumed by the first indefeasible one (that is it must be strictly more specific in terms of the type system). An associative order-independent operation of typed default unification is defined over TDFSs. This operation can be used to implement default inheritance in a type hierarchy and also more generally allow any two TDFSs with compatible types to be unified. Specificity of defaults is defined in terms of the type hierarchy, so a default value associated with a more specific (sub) feature structure overrides the default value associated with the less specific one. This requires that TDFSs record the type of the (sub) feature structure which introduced a default value into a given TDFS in order to avoid ordering effects from a series of unifications. Nixon diamonds cannot arise in the default inheritance system based on the type hierarchy as only orthogonal inheritance is allowed. However, unification of two conflicting default values from TDFSs of the same type will result in the least specific default value (i.e. $\top$); that is, the system is sceptical.

The underlying logic and formal details of TDFSs and the associated operation of typed default unification is given in Lascarides *et al* (1994, in press) and in Lascarides & Copestake (1995). Here we give some informal examples and motivation for these extensions to the formal framework proposed by Carpenter (1992). TDFSs can be abbreviated by collapsing identical parts of the double feature structures and marking the distinct default components by a slash, as in Figure 2.1 in which the reentrancy between path F:G and H is default as is the path consisting of the attribute H. We ignore tails in what follows since they are identical to the default component of the TDFS in the initial description.

Default inheritance in the TDFS system can be used to capture the fact that many lexical hyponymy relations are not absolute; for example, the telic role of a *book* is 'read' and *novel*, *monograph* and *dictionary* might all be defined as subtypes of *book*. However, in the case of *dictionary* the telic role will need to be overridden to 'refer-to'. Furthermore, as we shall see below there are reasons for thinking that the specification of telic roles as well as other aspects of lexical semantic specification must remain default 'beyond the lexicon'. The approach to defaults exemplified by the TDFS scheme, in contrast to most other approaches to default unification, allows exactly this by explicitly marking default specifications and ensuring that this marking persists through the default unification operation.

In addition, to inheritance via subtyping the TDFS lexical framework incorporates lexical rules which generate derived lexical entries from basic lexical entries represented as TDFSs. Most lexical rules are specified as type-to-type mappings without further constraints on the content of the input and output TDFSs beyond those specified in the relevant type declarations. In order to constrain the application of lexical rules to account for exceptions but allow for novel as well as conventionalised application (that is, deal with semi-productivity), a probabilistic interpretation is place on TDFSs which means that lexical rules may generate 'neological' unattested entries with very low probabilities associated with them. Further details are provided in Briscoe & Copestake (1995) where the approach is exemplified with reference to verbal diathesis alternations.

2.2 Systematic Polysemy

Any approach to lexical semantics which wishes to break away from the inadequate and unrevealing sense enumeration paradigm (Pustejovsky, 1995:39f) needs to distinguish systematic polysemy from idiosyncratic polysemy and homonymy. For example, the polysemy of *book* between the physical object and what it represents seems to be common to all words for representational artifacts not only including types of book but also films, pictures, and so forth. It is this type of lexical semantic ambiguity which a more adequate account of the lexicon should express generatively. Copestake & Briscoe (1995) and Copestake (1995) argue that systematic polysemy can be a result of constructional polysemy or semi-productive sense extension elaborating Cruse (1986) and others' informal distinction between sense modification and sense change.

2.2.1 Constructional Polysemy

Constructional polysemy is a consequence of sense modification in context where modification can be a consequence of further specification of an underspecified lexical entry, broadening of a specific entry with a default component, or coercion of quale to asserted status through semantic selection by predicates. (2a) represents a case of underspecification where the telic role of *reel* is '(un)wind' but the nature of the material to be (un)wound is unspecified and acquired from context.

(2) a The film/fishing/cotton reel broke.

b The cloud (of flies) was dense.

c John enjoyed the book.

(2b) represents a case of broadening where *cloud* has a default constituency of water vapour, but this can be overridden by an *of* complement. (2c) is a case of type coercion of *book* by the predicate *enjoy*, which selects for an eventive complement (and, by default, the one specified by the telic role – in this case 'reading'). The availability of both senses of *book* in the case of type coercion is illustrated by the possibility of what Pustejovsky (1995) terms co-predication – *John picked up and finished his beer* – supporting a treatment in which both senses are accessible from a single lexical entry. However, our current account does not deal with cases of coercion which involve non-constituent coordination, such as *John enjoys films and mending antique clocks* as a consequence of our decision to treat the coercion as internal to *enjoy* in order to deal with co-predication.

We argue that in these case there is no question of semi-productivity or semi-regularity. Rather if the appropriate linguistic context occurs the sense is modified in entirely predictable ways. We provide a fully formal account of each of these processes in the TDFS framework which serve to flesh out and formalise a number of mechanisms and phenomena which Pustejovsky (1995) discusses under the rubrics of co-composition, co-predication and type coercion. The regularity of the polysemy in type coercion of course only extends as far as the eventive interpretation of the bare NP. The precise event is open to further contextual specification as in *Judging by the number of pages missing, my pet goat obviously enjoyed that book.* I return to this below.

2.2.2 Semi-Regular Sense Extension

Semi-regular sense extension can be distinguished from constructional polysemy on the basis that it results in generation of a variant lexical entry by lexical rule and such rules are not fully productive. For example, the various conventionalised subcases of nominal 'grinding' which yield mass nouns denoting a substance originating from the count noun sense, such as wood-grinding (*oak tree*, *oak table*), and food-grinding (*a haddock*, *haddock filet* are quite regular within the appropriate lexical classes. However, there are exceptions caused by blocking (*cow/beef*, *pig/pork*), that is, preemption by synonymy. Furthermore, the general process of grinding, though applicable in the limit to any count noun denoting an individuated entity has very variable acceptability/conventionality, and conventionalised subcases differ cross-linguistically, suggesting the need for a linguistic overlaying of an ultimately more general conceptual transfer. An indication that grinding involves sense change rather than sense modification comes from the relative unacceptability of co-predication as compared to the book/contents polysemy discussed above: *?Sam fed and carved his lamb*, *Sam picked up and enjoyed that novel.*

Similar arguments can be made concerning the treatment of diathesis alternations: these also involve fairly regular sense extensions but are subject to lexical gaps which are not entirely predictable by restricting the application of lexical rules to semantic classes, even if these classes are narrowly specified (e.g. Pinker, 1989). Briscoe & Copestake (1995) show that verbs which undergo the dative alternation cannot be accurately specified in terms of narrow semantic classes, though a broad class specification is certainly necessary to restrict application of the rule. Nor does it seem likely that exceptions can be accounted for entirely in terms of blocking. Instead we propose that ultimately, acceptability is judged in terms of attested frequency and that low frequency lexical entries are utilised only in neologisms, for rhetorical effect or whatever.

The account of lexical rules in the TDFS framework captures semi-productivity by incorporating conditional probabilities that word forms will be associated with lexical entries into the control component of lexical rule application and providing a means for computing the overall probability of an interpretation. Probabilities are not a component of the TDFS LKRL but are more appropriately seen as part of the performance component of a language processing system. Thus, the acceptability of *He designed Bill a new house* as opposed to *He created Bill a new house* in American English is related simply to the non-occurrence of *create* in this construction in the experience of American English speakers and not to blocking or semantic differences between the two predicates. On the other hand neological use of *create* in this construction is interpretable because the lexical rule of dative can apply to *create* even though the resultant lexical entry is marked as unattested and of consequent low probability.

Blocking of sense extensions is not absolute but rather carries additional implicatures when it is overridden, as in *The drunken revellers were all too eager to tear into the roast cow.* In this sense blocking is a barrier to the establishment or institutionalisation of a sense rather than to its use. Briscoe *et al* (1995) provide an account of blocking which captures the pragmatic markedness of unblocking. The probabilistic account to some extent subsumes this within a more computationally tractable framework, but fails to deal with the implicatures generated by unblocking.

2.3 Semantic/Pragmatic Interactions

The framework of lexical semantic description we have developed is intended to be compatible with extant approaches to (computational) compositional semantics. Lexical semantic representations are couched in an event-based version of a default conditional variant of the predicate calculus, which is itself described using TDFSs. Lascarides (1995) and Lascarides *et al* (1996) following Frank & Reyle (1995) interpret these TDFS descriptions within an extended version of discourse representation theory integrated with an approach to discourse interpretation based on a conditional default logic (SDRT/DICE, Lascarides & Asher, 1993). Thus, the process of compositional semantic interpretation can be implemented using successive applications of typed default unification, and this allows us to formalise underspecification, broadening, coercion, co-composition and co-predication as an interaction between lexical semantic specification and compositional semantic interpretation.

The semantic part of a TDFS may contain default specifications which have persisted from the lexicon through compositional semantics. For example, *John enjoyed the book* will have a logical representation like $\exists(e, e', x, y) \wedge\ john(x) \wedge\ enjoy(e, x, e') \wedge\ book(y) \wedge\ P(e', x, y) \wedge\ *read(e', x, y)$ where P is a predicate variable over event(ualitie)s and **read* states that this is a reading event by default. For an example like *My goat enjoyed that book* this lexically-specified default is pragmatically inappropriate, so the interface to discourse interpretation must allow for such defaults to be overridden pragmatically. Lascarides (1995) gives a DRT-based semantics for the asterisk notation and shows how it can be integrated into the DICE framework in terms of two general principles: Defaults Survive and Discourse Wins. These two principles ensure that in the absence of conflicting discourse/pragmatic information lexically-specified defaults survive in the logical representation of the meaning of the input produced during parsing. However, in the presence of such conflicting information, for example that goats don't read, the dsicourse component both overrides the lexically-specified default and can supply a more appropriate refinement of the event variable in terms of more specific pragmatic information, for example, that goats eat indiscriminately. Briscoe *et al* (1990) provide some empirical evidence in support of this account of the lexical semantic / discourse processing interface, by demonstrating that logical metonymy is utilised only a) where the lexically-specified default interpretation is correct but the context may be uninformative or b) where the local context clearly overrides the lexically-specified default. Otherwise a non-metonymic complement is utilised. Less modular accounts of this interface (e.g. Hobbs *et al*, 1990) cannot explain the subset of a) cases in which the context is neutral other than by incorporating a lexical default into the pragmatic component. Evidence that these really are lexical defaults relating to word senses rather than concepts *per se* comes from examples like *?John enjoyed the doorstop* where although *doorstop* denotes a book, the telic role for the this lexical item does not provide a coherent refinement of the metonymic event leading to oddity.

Lascarides *et al* (1996) extend this account of the lexical semantic / discourse processing interface to account for coherence effects in co-predication and to provide a more adequate account of zeugma. The absence of crossed readings in examples like *John has a file and so does Bill* and the zeugmatic effects in ones like *John banked the money and then the plane* are usually explained in terms of the 'inaccessibility' of multiple lexical entries (senses) in coordination constructions. However, examples like *This chicken is corn-fed and delicious* or *This thesis has thousands of pages and is unreadable* which are not zeugmatic but involve polysemic senses suggest that this account is incomplete at best. In our framework the related senses of *thesis* are a consequence of constructional polysemy and are represented within one entry, but the senses of *chicken* result from the lexical rule of meat-grinding, suggesting that the former but not the latter should be zeugmatic given our theory combined with the standard account. However, it is also possible to construct zeugmatic examples using the same senses of *thesis* and *chicken* – *This chicken is beheaded and delicious*, *This thesis is excellent and on fire* – suggesting that the standard account is not capturing the relevant constraint. The causal relevance of the two conjuncts is critical to the discourse coherence of these coordinations. DICE captures these conditions in terms of conditions on the inference of the discourse relation *parallel* which is triggered by clausal coordination. It is shown that the interaction of polysemic senses with these conditions is what blocks crossed readings and cuases zeugma.

3 Conclusions

The approach that we have developed to lexical semantics and the organisation of the lexicon generally contra the proposals of Hudson (1986 (1995)) and of most computational systems (e.g. Hobbs *et al*, 1990) is modular in that we argue for a restricted LKRL with well-defined and restricted interfaces to syntax, compositional semantics and pragmatics. This more modular account others several theoretical and computational advantages, resulting in a stronger, more falsifiable theory which can be couched in a computationally-tractable formalism. Further research will be needed to decide whether the framework is empirically adequate. A number of immediate gaps and inadequacies can be identified: the interface between TDFS and DICE has not been implemented, the account of (un)blocking requires further work, and some types of non-constituent coordination involving type coercion remain outside the capabilities of the current proposals.

References

Bauer, L. 1983. *English word-formation.* Cambridge University Press, Cambridge, England.

Briscoe, E.J. and A. Copestake (1995) *Lexical rules in the TDFS framework,* Acquilex-II Working Papers, ??.

Briscoe, E.J., A. Copestake and B.K. Boguraev (1990) 'Enjoy the paper: Lexical semantics via lexicology', *Proceedings of the 13th Int. Conf. on Comp. Ling. (COLING-90),* Helsinki, pp. 42–47.

Briscoe, E.J., A. Copestake and A. Lascarides (1995) 'Blocking' in P. St. Dizier, P. and E. Viegas (eds.), *Computational Lexical Semantics,* Cambridge University Press, Cambridge, England.

Calcagno, M. (Barcelona) *Interpreting lexical rules,* Proc. of the Conference on Formal Grammar. 1995

Copestake, A. and E.J. Briscoe (1995) 'Semi-Productive Polysemy and Sense Extension', *Journal of Semantics, vol.12,* 15–67.

Carpenter, R. 1992. *The logic of typed feature structures.* Cambridge University Press, Cambridge, England.

Carter, R. (1976) 'Some constraints on possible words', *Semantikos, vol.1,* 27–66.

Copestake, A (1995) *Representing Lexical Polysemy, Proceedings of the AAAI Spring Symposium on Lexical Semantics*, Stanford CA.

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

Dowty, D. 1979. *Word Meaning in Montague Grammar.* Reidel, Dordrecht.

Dowty, D. (1989) 'On the semantic content of the notion 'Thematic Role'.' in G. Chierchia, B. Partee and R. Turner (eds.), *Property Theory, Type Theory and Natural Language Semantics,* D. Reidel, Dordrecht, The Netherlands, pp. 69–129.

Evans, R. and G. Gazdar (1989a) 'Inference in DATR', *Proceedings of the 4th Conference of the European Chapter of the Association for Computational Linguistics (EACL-1989),* Manchester, England, pp. 66–71.

Evans, R. and G. Gazdar (1989b) 'The semantics of DATR' in A. G. Cohn (eds.), *Proceedings of the Seventh Conference of the Society for the Study of Artificial Intelligence and Simulation of Behavior (AISB),* Pitman/Morgan Kaufmann, London, pp. 79–87.

Fillmore, C. (1968) 'The case for case' in Universals in Linguistic Theory (eds.), *E. Bach and Harms, T.,* Holt, Rinehart, Williams, New York, pp. 1–88.

Frank, A and U. Reyle (1995) 'Principle-based semantics for HPSG', *Proceedings of the 7th Meeting of the Eur. Assoc. for Comp. Linguistics,* Dublin.

Gazdar, G., E. Klein, G. Pullum & I. Sag 1985. *Generalized Phrase Structure Grammar.* Blackwell, Oxford, England.

Hobbs, J.R., M. Stickel, D. Appelt, and P. Martin (1990) *Interpretation as abduction,* SRI AI Center TR-499.

Hudson, R. (1986 (1995)) 'Identifying the linguistic foundations for lexical research and dictionary design' in D. Walker, Zampolli, A. and Calzolari, N. (eds.), *Automating the Lexicon: Research and practice in a Multilingual Environment,* Oxford University Press, Oxford, pp. 21–52.

Lascarides, A (1995) *The Pragmatics of Word Meaning, Proceedings of the AAAI Spring Symposium on Lexical Semantics*, Stanford CA.

Lascarides, A. and N. Asher (1993) 'Temporal Interpretation, Discourse Relations and Common Sense Entailment', *Linguistics and Philosophy, vol.16,* 437–493.

Lascarides, A., E.J. Briscoe, N. Asher and A. Copestake 1994, in press. Persistent associative default unification. *Linguistics and Philosophy.*

Lascarides, A. and A. Copestake (1995) *A logic for order-independent typed default unification,* Acquilex-II Working paper 60.

Lascarides, A., A. Copestake and E.J. Briscoe 1996, in press. Ambiguity and Coherence. *Journal of Semantics.*

Levin, B. 1993. *Towards a Lexical Organisation of English Verbs.* University of Chicago Press, Chicago.

Levin, B. (1985 (1995)) 'Approaches to lexical semantic representation (Lexical semantics in review: an introduction)' in D. Walker, Zampolli, A. and Calzolari, N. (eds.), *Automating the Lexicon: Research and practice in a Multilingual Environment,* Oxford University Press, Oxford, pp. 53–92.

Moens. M. and M. Steedman (1988) 'Temporal ontology and temporal reference', *Computational Linguistics, vol.14,* 15–28.

Moravsik, J. (1975) 'Aitia as generative factor in Aristotle's philosophy', *Dialogue, vol.14,* 622–636.

Pinker, S 1989. *Learnability and Cognition: The Acquisition of Argument Structure.* MIT Press, Cambridge, MA.

Pollard, C. and I. Sag 1987. *An information-based approach to syntax and semantics: Volume 1 fundamentals.* CSLI Lecture Notes 13, Stanford CA.

Pustejovsky, J. (1991) 'The Generative Lexicon', *Computational Linguistics, vol.17.4,* 409–441.

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

Pustejovsky, J. and B. Boguraev (1993) 'Lexical knowledge representation and natural language processing', *Artificial Intelligence, vol.63,* 193–223.

Pustejovsky, J. and F. Busa (1995) 'Unaccusativity and event composition' in Temporal Reference: Aspect and Actionality (eds.), *P. Bertinetto, V. Binachi, J. Higginbotham, M. Squartini,* Rosenberg and Sellier, Turin.

Riehemann, S. (1993) *Word formation in lexical type hierarchies,* M.Phil thesis, University of Tübingen, Germany.

Zaenen, A. (1993) 'Unaccusativity in Dutch: the interface of syntax and lexical semantics' in Semantics and the Lexicon (eds.), *J. Pustejovsky,* Kluwer, Dordrecht.