Enjoy the Paper: Lexical Semantics via Lexicology

Ted Briscoe & Ann Copestake
Computer Laboratory, Cambridge University
Pembroke Street, Cambridge, CB2 3QG, UK

Bran Boguraev
IBM Thomas J. Watson Research Center
PO Box 704, Yorktown Heights, New York 10598, USA

Abstract: Current research being undertaken at both Cambridge and IBM is aimed at the construction of substantial lexicons containing lexical semantic information capable of use in automated natural language processing (NLP) applications. This work extends previous research on the semi-automatic extraction of lexical information from machine-readable versions of conventional dictionaries (MRDs) (see e.g. the papers and references in Boguraev & Briscoe, 1989; Walker et al., 1988). The motivation for this and previous research using MRDs is that entirely manual development of lexicons for practical NLP applications is infeasible, given the labour-intensive nature of lexicography (e.g. Atkins, 1988) and the resources likely to be allocated to NLP in the foreseeable future. In this paper, we motivate a particular approach to lexical semantics, briefly demonstrate its computational tractability, and explore the possibility of extracting the lexical information this approach requires from MRDs and, to some extent, textual corpora.

1. Lexical Semantics

A theory of lexical semantics should provide an efficient representation of lexical semantic information in the paradigmatic plane which is capable of integrating with a genuinely compositional semantic account in the syntagmatic plane. Our starting point for this research is the work of Levin (e.g. 1985) and others on verbal alternations (diathesis), Pustejovsky (e.g. 1989) on lexical coercion and qualia theory, and Evans & Gazdar (e.g. 1989) on default inheritance within unification-based formalisms. It can be seen as a further contribution to the use of unification-based formalisms in linguistic description and specifically as an enriching of the minimal sort-based lexical semantic taxonomy incorporated into the Esprit ACORD system (Moens et al., 1989) and the SRI (Cambridge) CLE system (Alshawi et al., 1989). We propose a system in which a standard graph-based unification formalism, such as PATR-II, is augmented with minimal disjunction (of atomic terms) and minimal default inheritance (allowing only 'orthogonal' multiple inheritance in a manner similar to Evans & Gazdar's DATR). Using such a system we are able to see the beginnings of solutions to three problems concerning the integration of lexical semantics with a general theory of linguistic description and processing – alternations, coercion, and decomposition / representation.

The first problem emerges with systems, such as the Alvey Tools grammar (Carroll & Grover, 1989), which attempt to characterise the grammatical behaviour of lexical items in terms of sets of subcategorisation frames. Intuitively, this often seems arbitrary and inelegant because the occurrence of alternation seems to be semantically motivated. This problem has been discussed in connection with verbs mostly, but also arises with nouns and adjectives. For instance, in the Tools lexicon the verb *believe* has eight entries. Six of these separate entries relate to the same or a very similar sense of *believe*; namely, *believe3* (*Longman Dictionary of Contemporary English*, LDOCE) 'to hold as an opinion; suppose' which is a

relation between an individual (the believer) and a proposition (what is believed). Treating the various grammatical realisations of this sense of *believe* separately predicts that it is pure accident that they share the same sense. It also suggests that the range of possible alternations is unpredictable and must simply be listed from verb to verb. Most of the work on alternations has concentrated on attempts to characterise semantic classes of verbs which undergo similar alternations (e.g. Levin, 1985). This enterprise has not been particularly successful (Boguraev & Briscoe, 1989b), but in any case ignores or simply assumes the prior point that it is possible to construct a system in which there is just one entry for *believe3*.

Nevertheless, it seems correct that examples like *John believed that Mary was clever / Mary (to be) clever / Mary / the rumour* should be related to one entry for *believe* because this would allow us to account for the interpretation of *John believed Mary* as something like 'John believed something(s) that Mary asserted'; that is, as standing for some 'understood' proposition involving Mary. Pustejovsky (1989b) refers to this process as coercion and compares it to examples such as *John considers Mary a genius* where it is usual (e.g. in GPSG, Gazdar et al., 1985) to claim that *a genius* functions predicatively because the subcategorisation frame for *consider* forces this interpretation. In general, coercion is a problem in theories which take the syntactic aspect of grammatical realisation as primary, but would be a natural consequence of a theory which took the sense and the fact that *believe3* is a relation between an individual and a proposition as basic. In such an account an NP complement of a verb denoting a relation between an individual and a proposition would either denote a proposition 'directly' (*the rumour*) or be coerced to the appropriate semantic type (*Mary*).

When coercion occurs some additional information is required to 'flesh out' the elevated semantic type of the complement. Pustejovsky (1989) dubs this logical metonymy. In the case of *believed Mary* this is that it is some assertion of Mary's which is believed. This information appears to be inherited from the verb. In other cases, such as *John enjoyed (watching) the film*, *John began (reading) the book*, or *John finished (drinking) the beer*, it is more plausible that the missing information is provided by the lexical specification of the NP complements (cf: *John enjoyed (drinking) the beer*, *John finished (reading) the book*). Pustejovsky (1989, 1989b) and Pustejovsky & Anick (1988) propose that the lexical representation of nouns is enriched to include a specification of processes typically associated with the objects they denote and that, in cases of coercion, this information is utilised. In their terms, this is the telic role of the qualia structure of the noun.

We see the inheritance of this information from the verb or complement as a default process which operates in the absence of more marked pragmatic information. For example, one would normally enjoy (watching) the play, but it would not be difficult to construct a discourse context in which someone (say lecturer or student) enjoyed

(reading) the play, and so forth. So we propose that *enjoy* in this sense is a relation between an individual and an event and that, by default, nouns such as *film* or *play* inherit 'watch' as a specification of the typical event (process) in which they participate. The entry for *enjoy* will, also by default, state that in cases of coercion the specification of the process will be inherited from the nominal complement. In cases where the default is overridden by pragmatic information more specific instances of the entries for *enjoy* and/or *play* are created in which the defaults are replaced by pragmatically appropriate specifications. (The precise nature of the processes which trigger this or the retrieval of the relevant information we take to be a part of 'pragmatics' and not lexical semantics.)

One final (mostly methodological) point is that the approach we are advocating provides a slightly different viewpoint on the problem of lexical decomposition / representation. Early approaches to lexical semantics within the generative tradition were criticised for the arbitrariness of the representations produced. Following Dowty (1979), Pustojevsky (1989) and others, we suggest that one strategy for uncovering the optimal lexical representation, or level of 'decomposition', is to posit representations which provide elegant accounts of the interaction of lexical semantics with grammatical realisation and with compositional semantics. Pursuing this methodology, we have been led to a model of lexical semantic representation which supports a (somewhat enriched) compositional account of (sentence) meaning by enriching lexical representations of nouns and collapsing those for verbs with alternate grammatical realisations. In this framework, there are still many inferences which are not captured at the level of lexical organisation, but we argue that these inferences are 'pragmatic' in the sense that they are not based on default processes operating within the lexicon. Thus, our position is opposed to that of Hobbs et al. (1987) who argue that there is no distinction between lexical semantics and general knowledge. In our approach, simple default 'lexical' inference procedures do quite a lot of work. Of course, the way is always open to us to argue that any inference which cannot be captured by these procedures is 'non-default'. Nevertheless, in section 3 we argue that this distinction is supported by natural data both in terms of the rarity of non-default cases and also the marked, informationally-rich nature of the contexts in which lexical defaults are overridden. Thus, our approach gives us a handle on which aspects of lexical meaning should be represented in the lexicon and therefore on the type of information we want to extract from our MRDs.

Figure 1

2. An Implementation

It is possible to implement a system capable of coercion and default specification using a unification-based formalism extended with 'orthogonal' default inheritance of (paradigmatic) lexical specifications. We also make use of minimal disjunctive specifications to allow for the range of grammatical alternation within one sense of a predicate. Our prototype extends PATR-II (Shieber, 1986) with disjunction of atomic terms and uses the template mechanism to impose a natural subsumption ordering on the lexical taxonomy which defines the inheritance network. The taxonomy implicit in the fragment implemented so far is shown in Figure 1. This taxonomy is adequate to cover the metonymies discussed in this paper and others discussed in Pustojevsky (1989). (Numbers on concepts are related to LDOCE sense numbers.)

An entry for *book* is given in template form in Figure 2a. Its position in the network in Figure 1 defines the pattern of inheritance for the qualia structure.

Lexical entry for "book":

book - 1 N Literature PhysObj;

Dag for "the book":

```

[CAT:NP
 SEMF'S:[CAT:OBJ
          TYPESHIFTED: FALSE]
 TRANS:[DET:DEFINITE
         PRED:BOOK1
         VAR: <DAG61> = REF'25
         ARG1: <DAG61>]
 QUALIA:
   [TSTRUCT:
     [TRANS:[PRED:?READ
              VAR: <DAG62> =[]
              EVENT: <DAG62>
              ARG1: <DAG63> =[]
              ARG2: <DAG61>]
     COMBINES:[ FIRST:
                  [TRANS:[VAR:<DAG63>]]]]]
  
```

Figure 2a

The telic role for *book* is thus inherited from the default role associated with 'Literature'. The entry will also inherit information from 'PhysObj' but the orthogonality constraint rules out conflicts with the attributes inherited from 'Literature'. In fact the template 'PhysObj' does not contain any information about the telic part of the qualia structure.

The DAG for the NP *the book* is also shown in Figure 2a. This still denotes an object; when combined with a normal, non-coercing verb the telic role makes no contribution to the semantic structure. However some grammar rules allow type-shifting; one allows NPs with an associated telic role to be type-shifted to be equivalent to untensed VPs and to denote events.

Figure 2b shows the NP after application of this rule. Once type-shifted, the logical formula associated with *the book* is the same as that associated with *reading the book*, except that the question mark indicates defeasibility and could be interpreted as 'possibly(P = read) & P(e' j x)'.

```

[CAT:NP
 SEMFS: [CAT:EVENTUALITY
  TYPESHIFTED:TELIC]]
TRANS: [PRED:AND
  VAR:<DAG39>=REF26
  ARG1: [DET:DEFINITE
    PRED:BOOK1
    VAR: <DAG40>=REF25
    ARG1:<DAG40>]
  ARG2: [PRED:READ
    VAR:<DAG39>
    EVENT:<DAG39>
    ARG1:<DAG41>=[]
    ARG2:<DAG40>]]
COMBINES: [FIRST:
  [TRANS: [ VAR:<DAG41>]]]]

```

Figure 2b

In (1a) we show the formula which can be read off the DAG in Figure 2b given straightforward assumptions about the semantic interpretation of the formalism (e.g. Moore, 1989). The lexical entry for *enjoy* specifies that its complement must denote an event which can be syntactically an NP or progressive VP and that, if the NP is type-shifted, the telic role supplies the understood predicate. The resulting formulae associated with the VP and S are shown in (1b,c).

- (1)
- a) $\lambda x \lambda e' \exists y ?\text{read}(e' x y) \ \& \ \text{book}(y)$
 - b) $\lambda x \lambda e' \exists e' y \text{past}(e) \ \& \ \text{enjoy}(e x e') \ \& \ ?\text{read}(e' x y) \ \& \ \text{book}(y)$
 - c) $\exists e e' y \text{past}(e) \ \& \ \text{enjoy}(e j e') \ \& \ ?\text{read}(e' j y) \ \& \ \text{book}(y)$

We follow Hobbs (1985), Alshaw et al. (1989), Moens et al. (1989) and others in using an event-based calculus for reasons of computational tractability, and also because distinctions amongst types of events are likely to be important in the characterisation of the recovery of understood predicates in logical metonymies. In a fuller account it would be possible to constrain the type of event selected by a particular verb; for instance, *enjoy* might be constrained to unify by default with the telic role of a noun if this specified a process or culminating event. This would predict the relative oddity of examples such as *John enjoys his house*, in which we assume that the telic role is something like 'living in' and that this specifies a state rather than process. It would also be possible to alter the aspect of qualia structure selected by a particular verb. An example like *John regrets that book* by default receives an interpretation in which 'writing' is selected to flesh out the metonymy. In this case, we might specify that *regret*, in contrast to *enjoy*, selects the agentive path in the noun's qualia structure.

Another area in which this approach to lexical semantics is suggestive relates to adjectival modification. It is well-known that adjectives such as *good*, *fast*, *long*, and so forth, have meanings which are hard to specify independently of some 'aspect' of the noun they modify. Pustejovsky (1989) suggests that in examples like *fast car*, *fast typist*, or *fast waltz*, *fast* should be treated as a modifier of the telic role associated with these nouns, so that these examples can be paraphrased *fast car to drive* or *fast waltz to dance*. The adjective *long* appears to be (at least) ambiguous between a telic role modifier and a formal role modifier – *a long book* can either be a comment on shape, size or number of pages, or a comment

on the length of time required for reading. In the event-based calculus we adopt we could associate the logical form in (2b) with the interpretation of (2a) where *long* is a telic role modifier.

- (2)
- a) John enjoyed the long book
 - b) $\exists e e' e'' x y \text{enjoy}(e j e') \ \& \ ?\text{read}(e' j y) \ \& \ \text{book}(y) \ \& \ \text{long}(e'') \ \& \ ?\text{read}(e'' x y)$

However, note that it would be inappropriate to automatically conflate the events e' and e'' because this would predict that John's reading of the long book was necessarily a long event which, whilst plausible, is not entailed under this interpretation of *long*. In order to avoid this effect using unification-based techniques it is necessary to explicitly copy the structure that specifies the telic role.

We suggested in section 1 that NPs, such as *the fact*, can denote propositions 'directly'. Similarly, we think that there is no metonymy involved in examples such as *John enjoyed the experience / film-making* and so forth. In these cases, we claim that the NPs in question denote events 'directly'. Thus, we are lead to an 'ontologically promiscuous' semantics (Hobbs, 1985). However, recent developments in model-theoretic semantics which treat properties as basic entities (e.g. Chierchia & Turner, 1988) support this position. Indeed the interpretation of event-denoting NPs in complement position with *enjoy* strongly suggests that these NPs must be analysed as denoting propositional functions since their 'missing argument' must be associated with the subject of *enjoy*. For instance, *John likes marriage* can mean that John likes the institution but *John enjoys marriage* can only mean that he enjoys being in the state of marriage (to someone).

3. Data concerning Logical Metonymies

The previous sections have demonstrated the nature of the phenomenon of logical metonymy and have outlined a computationally-tractable unification-based treatment. A crucial aspect of this treatment is that, with the predicates we have considered, the missing information is supplied, by default, by the qualia structure of the head noun in the type-shifted complement. In order to demonstrate the presence of logical metonymies in naturally-occurring text and to evaluate the plausibility of our default approach, we examined data drawn from the Lancaster-Oslo/Bergen (LOB) corpus containing predicates capable, in principle, of coercing the type of their complements.

A set of type-coercing predicates similar to *enjoy* was obtained by extracting verbs coded to take both NP and progressive or infinitive VP complements in LDOCE (see Boguraev & Briscoe, 1989b for an account of these codes and the extraction techniques). Further manual editing of this list led to 24 predicates which we felt were capable of exhibiting logical metonymies parallel to that of *enjoy*. To date, we have analysed all the data obtainable from the LOB corpus for seven of these predicates. The results of this analysis are summarised in Figure 3. (Numbers after predicates refer to LDOCE sense numbers.)

Pred	Prog	Inf	NP	Ev	Met	Prag
enjoy1	6	/	59	21	25	4
prefer1	4	30	30	10	13	1
finish1	8	/	31	8	23	6
start1,3	45	28	63	42	21	0
begin1	1	57	11	8	3	2
miss5	3	/	24	10	13	4
regret1	2	/	17	14	0	0

Figure 3

Columns headed NP, Inf(imitive) and Prog(ressive) show the number of times each predicate occurred with this type of complement. (A stroke in these columns indicates that this complement type would be ungrammatical with a particular predicate.) The remaining columns give further information about the NP complements. Ev(ent) indicates the number of times that the NP complement was judged to denote an event (or in a few cases a proposition) directly. Met(onymic) indicates the number of times we judged that coercion had occurred. And Prag(matic) indicates the number of times that we judged the understood predicate was not recovered via the head noun's qualia structure in the metonymic cases. In some cases, the number of NP complements is greater than the sum of Event and Metonymic because we felt unable to classify some examples. These examples were either (semi) idiomatic, such as *miss the boat*, or involved NPs whose status was unclear because of modification of the head noun, such as *enjoy the warm evening*.

The first thing to note about Figure 3 is the comparatively high numbers of metonymic examples relative to the complete sets. It is instructive that the apparently more complex metonymic complement pattern is selected quite frequently despite the availability, with all these predicates, of an explicit VP complement pattern. For instance, *enjoy* and morphological variants occurs 65 times in the relevant sense and coerces its NP complement in 25 of these cases. The second and crucial observation, from the perspective of our default theory of the recovery of the understood predicate in the metonymic cases, is that the numbers in the Pragmatic column are relatively low by comparison with the total number of metonymic examples. Given the default theory, we would expect most metonymic examples to be resolvable via the head noun's qualia structure and there to be relatively few 'pragmatic' examples involving less constrained and more complex inferences, and, in fact, these cases represent about 17% of the metonymic examples and about 4% of the total set of examples considered. Further examination of these 17 cases revealed that, in most, the immediate context was informationally-rich and therefore marked enough for the appropriate pragmatic inference to go through. For example, compare the a) examples with the b) examples in (3).

(3)

- a) Willie enjoyed the hot sweet tea, standing on the deck in the cool of the night
- b) She can lie back and enjoy her baby until the midwife, knowing the afterbirth is ready to pop out, ...
- a) Loddon paid his own account, finished his cigarette, and got up.
- b) The book was never finished, for his illness and death intervened while he was in the course of writing it.

a) If you prefer a Burgundy try a 1955 Charnes Chambertin costing round 1 pound.

b) Then again so many people much prefer the sea or river to the baths. Having learned to swim in the sea ...

In each of the a) examples we think that the understood predicate is supplied via the qualia structure of the head noun in the NP complement. In the b) cases, it seems implausible that the telic role of babies is to be cuddled, or that seas or rivers are (mainly) for swimming in. The agentive role of book will specify the predicate 'write', so we could treat *finish* as selecting this role by default and this would, in fact, deal with four of the 'pragmatic' cases, but others would become 'pragmatic' since in our implementation only one unification path into the qualia structure can be selected by default. However, in each of these examples the context shown provides enough information to infer the relevant predicate. It is not always the case that the (remainder of the) context provides the relevant information or intuitively seems so 'rich' in the default cases.

Another way in which we can evaluate the default theory is by considering the status of the predicates which are supplied explicitly when a VP complement is selected. We might expect VPs to be selected precisely in those situations when defaults based on qualia structure would lead to the wrong interpretations. We tested this idea by examining the VP complements of *start*. In many cases, the predicates were intransitive, ditransitive, and so forth, so that the hypothesis did not apply. However, in the straightforwardly transitive cases 21 examples exhibited clear non-default predicates, such as *started to open the bottle*, *started to play a Waltz*, or *started flirting with the first pretty girl that you met*, whilst only 4 cases arguably involved default predicates recoverable from the head noun's qualia structure – *start making a fuss*, *started to fire distress rockets*, *started pulling the communication cord*, and *started making bubbling noises*.

This analysis is hardly conclusive, however it does, we think, demonstrate that logical metonymies occur quite regularly with certain predicates in natural text. We have also provided some evidence that default inference based on lexical organisation (in this case the qualia structure of nouns) would succeed in a large number of cases. Furthermore, there seems to be some support in this data for the claim that contexts in which 'pragmatic' recovery of the understood predicate occurs are quite informationally-rich and would therefore constrain an otherwise rather unconstrained process. Finally, we have shown that, in the case of *start* there is evidence that VP complementation is chosen when default recovery of understood predicates on the basis of qualia structure would lead to the wrong interpretation.

4. Acquiring Lexical Semantic Information

In this section, we describe three exploratory studies aimed at the (semi-)automatic acquisition of qualia structure, in particular telic roles, from MRDs. The first involves exploiting subject and box codes in the LDOCE MRD (see papers in Boguraev & Briscoe, 1989 for a full description), while the second is based on an analysis of the LDOCE definitions. These techniques are aimed at allowing qualia structure to be inherited appropriately; the third attempts to determine the predicates associated with a word by analysis of dictionary definitions and, to some extent, more general corpus material.

The machine-readable version of LDOCE contains some residual 'database-like' features which do not appear in the printed dictionary. These include a taxonomy of many words in terms of 'subject matter'. This taxonomy defines a 'flat' hierarchy of, at most, two levels and many relationships are left implicit; for instance, 'sports' is a main category with subdivisions such as 'archery' but 'football' is a main category with subdivisions such as 'rugby'. Nevertheless, this taxonomy can be used to identify 'lexical conceptual paradigms' (Pustejovsky & Anick, 1988); for example, there is a class 'beverages' (147 word senses), a class 'motion pictures' (113 word senses), and a class 'literature' (377 word senses). These words could straightforwardly be associated with the 'deeper' inheritance network given in Figure 1 with default telic (and possibly other) roles, such as 'drinking', 'watching' and 'reading' associated appropriately. There are a few problems though, for instance the category 'beverages' includes *publican*, and 'motion pictures' includes *usherette*. It is possible to exclude these examples from the target network by making use of box codes which, amongst other things, associate semantic features with nouns, because the exceptions are coded 'animate' and 'human'. Nevertheless, this approach is limited because the LDOCE semantic taxonomy will undoubtedly not contain all the classifications which eventually will prove desirable and there will be errors of omission in its construction. In addition, we are utilising an idiosyncratic feature of the LDOCE MRD, whilst we would like our extraction techniques to be generally applicable.

An improvement to this approach is to utilise taxonomies constructed from the dictionary definitions. For example we have built a taxonomy of substances by extracting the genus senses of LDOCE definitions in which 145 word senses such as *Burgundy* appear directly or indirectly under the main nominal sense of *drink*. We are currently investigating an approach whereby lexical entries inherit some of their structure from higher nodes in the taxonomy. Qualia structure could thus be inherited from word senses rather than abstract templates; for example *Burgundy* would inherit its telic role from the noun *drink*. If abstract templates were still needed they could be inserted into the inheritance hierarchy at the appropriate points.

The approaches above only specify how the qualia structure is inherited, rather than how it is initially determined. In recent work, the IBM lexical systems group have used their lexical database system (e.g. Neff & Boguraev, 1989) with a number of MRDs to generate lists of predicates which are applied to books by searching through definition fields for the occurrence of *book* in a position denoting 'typical object' of the headword. For instance, LDOCE defines *sag* with '(of a book, performance, etc.) to become uninteresting during part of the length'. Using these techniques with three dictionaries resulted in the following list of verbs: *abridge, abstract, annotate, appreciate, autograph, ban, bang about, borrow, bring out, burlesque, bowdlerize, call in, castigate, castrate, catalogue, censor, chuck away, churn out, classify, collate, commission, compile, consult, cross-index, dramatize, entitle, excoriate, expurgate, footnote, page, pirate*. It is obvious that this technique yields specific, often rare, predication with typical objects. Whilst qualia structure is likely to involve typical predications with specific (classes of) nouns.

In order to automatically obtain typical (frequent) predications of *book*, four corpora were searched –

LDOCE example sentences, the Brown corpus, 1.2 million words of *Readers' Digest*, and 26 million words of tapes from the American Publishing House for the Blind. Analysing those citations in which *book* occurs as direct object revealed that *read* and *write* are the two most common predicates across the four corpora, although there are considerable differences within each corpus (see Boguraev et al., 1989 for details). This approach could and should be extended in several ways, for instance by dealing with semantically related nouns such as *novel*, and, of course, by attempting a similar analysis for many more nouns. Nevertheless, these preliminary results do suggest that a noun's qualia structure should be recoverable from MRDs and corpora in a semi-automatic way.

5. Conclusion

We have attempted to motivate an approach to lexical semantics which enhances the representation of nouns in terms of their qualia structure. We have shown that incorporating this information into a default inheritance hierarchy and enriching the notion of compositionality to allow for type-shifting of NPs allows for a computationally tractable and plausible account of logical metonymy. We have, however, said very little about what qualia structure is. Whilst Pustejovsky (1989) relates this idea back to Aristotle's four causes, we think that for the purposes of the computational implementation described above we need only assume that qualia structure constitutes (part of) the lexical information associated with a word sense, in the sense that it is the information which is most accessible given the organisation of the inheritance network. One could imagine that other more general or 'encyclopedic' information concerning concepts would simply be less accessible or 'close' in terms of the same network. The preliminary work with MRDs/corpora suggests that both types are recoverable semi-automatically.

6. References

- Alshaw, H., Carter, D., van Eijck, J., Moore, R., Moran, D., Pereira, F., Pulman, S. & Smith, A. 1989. Final Report: Core Language Engine. SRI (Cambridge) Technical Report, Project No. 2989.
- Atkins, B.T. 1988. Course Notes. ESF Summer School on Automating the Lexicon. To appear in Atkins, B.T. & Zampolli, A. *Automating the Lexicon*, Oxford University Press, Oxford.
- Boguraev, B. & Briscoe, E. 1989. *Computational Lexicography for Natural Language Processing*. Longman/Wiley, London/New York.
- Boguraev, B. & Briscoe E. 1989b. Utilising the LDOCE grammar codes. In Boguraev & Briscoe 1989.
- Boguraev, B., Byrd, R., Klavans, J. & Neff, M. 1989. From structural analysis of lexical resources to semantics in a knowledge base. IBM Research, Mimeo.
- Carroll, J. & Grover, C. 1989. The derivation of a large computational lexicon of English from LDOCE. In Boguraev & Briscoe 1989.
- Chierchia, G. & Turner, R. 1988. Semantics and property theory. *Linguistics & Philosophy*, 11, 261-302.
- Dowty, D. 1979. *Word Meaning and Montague Grammar*. Reidel, Dordrecht.
- Evans, R. & Gazdar, G. 1989. Inference in DATR. *Proc. of 4th Eur. ACL.*, Manchester, pp.66-71.

- Gazdar, G., Klein, E., Pullum, G. & Sag, I. 1985. *Generalized Phrase Structure Grammar*. Blackwell, Oxford.
- Hobbs, J. 1985. Ontological promiscuity. *Proc of 23rd ACL.*, Chicago, pp.61-9.
- Hobbs, J., Croft, W., Davies, T., Edwards, D. & Laws, K. 1987. Commonsense metaphysics and lexical semantics. *Computational Linguistics*, 13, 241-50.
- Levin, B. 1985. Lexical Semantics in Review. MIT Working Papers on the Lexicon & in Walker et al. 1988.
- Moens, M., et al. 1989. Expressing generalisations in unification-based grammar formalisms. *Proc. of 4th Eur. ACL.*, Manchester, pp.174-81.
- Moore, R. 1989. Unification-based semantic interpretation. *Proc. of 27th ACL.*, Vancouver, pp.33-41.
- Neff, M. & Boguraev, B. 1989. Dictionaries, dictionary grammars and dictionary entry parsing. *Proc. of 27th ACL.*, Vancouver, pp.91-101.
- Pustejovsky J. 1989. Current Issues in Computational Lexical Semantics. *Proc. of 4th Eur. ACL.*, Manchester, pp.xvii-xxv.
- Pustejovsky, J. 1989b. Type Coercion and Selection. *Proc. of WCCFL VIII*, Vancouver.
- Pustejovsky, J. & Anick, P. 1988. On the semantic interpretation of nominals. *Coling88*, Budapest, pp.518-23.
- Shieber, S. 1986. *An Introduction to Unification-based Approaches to Grammar*. U. Chicago Press, Chicago.
- Walker, D., Zampolli, A. & Calzolari, N. 1988, in press. *Automating the Lexicon: Research and Practice in a Multilingual Environment*. Cambridge University Press, Cambridge.