



Peeling an Onion: The Lexicographer's Experience of Manual Sense-Tagging

RAMESH KRISHNAMURTHY^{1,*} and DIANE NICHOLLS²

¹*ramesh@easynet.co.uk*; ²*dianenicholls@email.msn.com*

Abstract. SENSEVAL set itself the task of evaluating automatic word sense disambiguation programs (see Kilgariff and Rosenzweig, this volume, for an overview of the framework and results). In order to do this, it was necessary to provide a 'gold standard' dataset of 'correct' answers. This paper will describe the lexicographic part of the process involved in creating that dataset. The primary objective was for a group of lexicographers to manually examine keywords in a large number of corpus contexts, and assign to each context a sense-tag for the keyword, taken from the Hector dictionary. Corpus contexts also had to be manually part-of-speech (POS) tagged. Various observations made and insights gained by the lexicographers during this process will be presented, including a critique of the resources and the methodology.

Key words: context, corpus, evaluation, lexicography, part-of-speech tagging, word sense disambiguation, sense-tagging

1. Introduction

Lexicography is a multi-faceted activity. Far from being a harmless drudge, a lexicographer needs to access a wide range of linguistic and cultural knowledge and employ analytical and editorial skills in a creative process that is neither wholly art nor wholly science. Using corpus contexts is a relatively recent methodology (Sinclair, 1987). It can add enormously to intuition and introspection, especially in terms of accuracy and frequency. It can also be expensive and time-consuming (not to mention repetitive and tedious for the lexicographer).

Getting any two human beings to agree on anything can be difficult, and lexicographers are generally more disputatious than average. In this particular task, knowing that a crucial aspect of our role was in providing independent taggings in order to gauge the degree of consensus among human taggers, the lexicographers deliberately did most of the work in isolation. We knew that others might be analysing the same word, but did not communicate with them about it in any detail.

Six highly experienced lexicographers participated in the manual tagging,¹ and the whole exercise spanned approximately two months. In late May, the lexicographers were sent draft tagging instructions, Hector dictionary entries (see Atkins, 1993), and 100 corpus contexts for the test word **promise**. This was followed by a face-to-face meeting in Brighton in early June, to compare experiences, fine-

tune the procedures and so on. Thereafter, there was very little communication, apart from the occasional email or telephone call. The deadline was fixed for 17th July. Subsequently, there was a brief second phase, during which disagreements between human taggers and/or Hector's tagging were reviewed by three of the lexicographers (see Kilgariff and Rosenzweig, this volume).

This paper is based on the experiences and comments of all the lexicographers who took part, but the responsibility for any errors or misrepresentations lies with the authors. Throughout this paper, Hector dictionary headwords and POS tags are in bold, sense-tags are underlined, context words are in capitals, and corpus contexts are in italics.

2. Procedures

The Hector dictionary entries consisted of headwords with numbered senses and subsenses, each associated with a mnemonic codeword, some clues (syntax, collocates), optionally a subject field or register specification, a definition, and one or more examples (often with additional notes). The corpus contexts were numbered, and the word to be tagged was the first instance of the headword in the last sentence of the context. Lexicographers were to return files containing: context number, Hector sense-mnemonic, and part-of-speech. Various options were available for complex cases, with unassignable, typo (i.e. typographic error) and proper noun as special tags.

Specific instructions to lexicographers included the following:

- (a) If a corpus instance clearly matches a subsense, assign the subsense. If the match is not clear, assign the main-level sense (e.g. sense 4.1 of **promise** is 'make a promise', so a corpus instance such as *He muttered a promise* should not be assigned to sense 4.1, but to the more general main sense 4).
- (b) Tag gerunds as **a** (for adjective) or **n** (for noun). Note that 'promising' may be a noun form, but it is not the common noun form of 'promise' (which is "promise"!), so it would be misleading to tag it **n**. In such cases, use the POS-tag "?".
- (c) Treat heads of reduced relative clauses (i.e. -ed and -ing forms) as verb occurrences.
- (d) When assigning POS, do not treat the lexical unit as something larger than the single word (even if it is linguistically accurate to do so). Give the POS for the target word alone.
- (e) In return files, the first column is for reference number, second for mnemonic, third for POS.
- (f) Where there is not enough context to be certain which sense applies, write 'no-context' in the fourth column.

- (g) Use disjunction ('mnemonic1 or mnemonic2 or mnemonic3') in the mnemonic column.
- (h) In general, use disjunction rather than opting for just one tag, or using the ? or x (for one-off 'exploitations') suffixes, where you are unsure which tag to apply.

3. Time Constraints and Working Methods

The average working rate was 66 contexts per hour (reported rates varied between 40 contexts per hour to over 100). The rates were lower at the beginning of the task, and also varied according to the difficulty of each word. All the lexicographers found that they worked faster as they became more accustomed to the dictionary sense-divisions and the mnemonic tags. Also, whereas they tended to look at the whole context initially, for later contexts a quick glance at the immediately surrounding context was often sufficient.

4. Hector Dictionary Entries

4.1. SENSE DIVISIONS

Sense-division is a notoriously difficult area of lexicography (Stock, 1984, Kilgarriff, 1998), and one that can give rise to heated and acrimonious debate on occasion. The lexicographers in this exercise were quite critical of the Hector sense divisions that they were being compelled to apply to the corpus contexts. They frequently suggested that the distinctions made in Hector were not sufficiently fine to reflect the corpus contexts:

accident: add an extra sense or sub-sense between sense 1 (crash: "an unfortunate or disastrous incident not caused deliberately; a mishap causing injury or damage...") and sense 2 (chance: "...something that happens without apparent or deliberate cause; a chance event or set of circumstances.") to cover broken windows and spilt coffee (rather than car crashes or nuclear meltdowns), characterised by contexts such as *have an accident*, *it was an accident*, etc.

accident: the expression 'accident and emergency' (used to denote a medical speciality and hospital department) should be treated as a separate sense.

generous: sense 3 kind (for definition, see Appendix to Kilgarriff and Rosenzweig, this volume) is really two different senses; the definition is in two halves ("recognizing positive aspects" and "favouring recipient rather than giver"); if subdivided, the second definition could then be expanded to cover *generous conditions of employment*, *generous odds*, etc.

hurdle: add a sub-sense, ‘threshold, qualifying standard’, to sense 5 (obstacle: “(in metaphorical use) an obstacle or difficulty”) for contexts like *the 5% hurdle in elections, the quality hurdle for TV franchises*.

knee: add a sub-sense to sense 1 (patella: “the joint between thigh and lower leg”) for ‘marker of height on the body’ (cf Hector dictionary 4th and 5th examples ... *any hemline above the knee* ... and ... *you’d be up to your knees in a bog*.)

shake: the physical, literal sense of ‘shake someone off’ or ‘shake someone’s hand/arm off’ is missing in Hector, but present in the corpus lines.

slight: important to split sense 1 (tiny: “very small in amount, quantity, or degree”) to distinguish examples with negative force (mainly predicative) from those with positive/neutral force (mainly attributive); (cf. ‘little’ and ‘a little’ etc).

Very few comments suggested that the Hector senses were too finely distinguished:

generous: often difficult from the context to decide between sense 3 (kind) and sense 1 (unstint) (for definitions, see Appendix to Kilgarriff and Rosenzweig, this volume), so create an umbrella sense covering ‘a person or their character’.

4.2. GRAMMAR

The Hector dictionary aimed to place semantics first, with syntax merely acting in a secondary, supporting role. This meant that syntactic coding could not be taken as definitive. Also, the coding did not distinguish obligatory from optional syntactic features. The lexicographers certainly noticed many instances of corpus contexts which matched a Hector sense in terms of meaning, but did not match the sense’s syntactic specification.

band: senses 1 (mus: “a group of musicians”) and 2 (group: “a group of people with a common interest, often of a criminal nature”) are labelled as **nc** (countable noun) in Hector, but need to be additionally labelled as collective nouns, because they can be used with a plural verb.

behaviour: sense 1 (socialn: “the way in which one conducts oneself”) is marked **nu** (uncountable noun), but there are several **nc** instances.

consume: ‘consumed with’ is not covered in the syntax or examples, yet is common.

invade: senses 2 (takeover: “(of persons, less commonly of animals/things) to make an incursion into an area, etc”), 2.1 (infest: “(of parasite/disease) to infest an organism”), and 2.2 (habitat: “(of animal/plant) to spread into a new area/habitat”)

are only marked **vt** (transitive verb) and need to be additionally marked **vi** (intransitive verb).

sanction: sense 2 (penalty: “a penalty for not complying with a rule, a coercive/punitive measure”) is marked **nc**, but there are **nu** contexts.

4.3. MULTI-WORD ITEMS

One problem raised was how to code an item, when an **n-mod** (noun used as modifier) sense was specified in the dictionary and the item was part of a compound, in cases where the whole compound was modifying, rather than the headword on its own (e.g. ACCIDENT in *personal accident insurance*). The variability of phrases was also a matter of concern:

bother: *if you can be bothered* and *couldn't be bothered* didn't exactly match the phrase “can't be bothered” in the entry for **bother**, yet they were clearly closely related to it.

4.4. DEFINITIONS

One consistent difficulty was with the distinctions made between animate and inanimate entities.

behaviour: a lot of contexts show institutions (e.g. banks and unions) acting as conscious entities.

The use of near-synonyms in separate definitions also caused problems. In the entry for **bitter**, sense 2 (feelings) says “(of people, their remarks, or feelings)”, and sense 4 (unpleasant) says “(of an experience, situation, event, emotion, etc.)”. The difference between ‘feelings’ and ‘emotion’ was difficult to resolve.

4.5. EXAMPLES

Occasional criticisms were made of the examples given in the Hector dictionary:

shake: *They were badly shaken by the affair* was tagged by the lexicographers as the verb sense disturb (“(of a person, event, phenomenon, etc) to disturb, disconcert, or upset the equilibrium (of a society, group, person)”) or as the adjective sense troubled (“(of a person) severely upset or shocked, as after an accident, bad news, etc”). The distinction is not clear in the Hector examples.

5. Lexicographers' Observations on the Corpus Contexts

Once the human taggers had established a working procedure, familiarized themselves with the various aspects of the Hector dictionary outlined above, received their individual assigned words and digested the sense definitions available to them, they then turned to the corpus contexts for each word. Although the majority of contexts were clear and simple to tag, the taggers encountered a number of difficulties.

5.1. INSUFFICIENT CONTEXT

Some contexts, particularly the more literary or specialised, were too brief for a sense to be assigned. Others were either too vague or the dictionary sense distinction didn't help. For example, in **bet (n)**, the senses are either wager:n (an arrangement between two or more people whereby each risks a sum of money or property on the outcome of a race, competition, game or other unpredictable event), or speculation:n (an act of speculation or an opinion about an outcome or future situation). These two have the same syntactic information and, semantically, only differ regarding whether money or property is involved. There were at least seven contexts where it was not clear whether money or a simple statement of certainty was involved, so the tagger could not know which of two possible senses to assign. For example:

700235 *Opinions are opinions, of course, but when they are so uniform and consistent (particularly about a polling result which can be interpreted completely differently), we readers have to ask whether you might collectively be trying to tell us something? TODAY a contest will begin that may finally settle a <>BET<> made 21 years ago.*

700296 *Temple Cowley Pool: No, I have not lost my <>BET<>!*

Some contexts simply made no sense at all to the tagger or at least left the taggers with a feeling that there was a large gap in their world knowledge, or a sense missing in the dictionary of which they themselves were unaware²:

700004 *In fact it is not all that obvious, and I did take the precaution of simulating it on a computer to check that intuition was right. Grudger does indeed turn out to be an evolutionarily stable strategy against sucker and cheat, in the sense that, in a population consisting largely of grudgers, neither cheat nor sucker will <>INVADE<>.*

700007 *The locally stable strategy in any particular part of the trench lines was not necessarily Tit for Tat itself. Tit for Tat is one of a family of nice, retaliatory but*

forgiving strategies, all of which are, if not technically stable, at least difficult to <>INVADE<> once they arise.

5.2. ENCYCLOPAEDIC OR 'REAL-WORLD' KNOWLEDGE

A broad bank of encyclopaedic or real-world knowledge and the ability to make assumptions and leaps of logic were a distinct advantage. The taggers could draw on their own experience of the world when assigning senses to contexts. This advantage was very much in evidence in the tagging of **band**, where tagging would have been difficult, if not impossible in some cases, if the tagger had not known, for example, that 'Top of The Pops' is a popular music programme on British TV:

700277 *'It'd be something they remembered.' 'It's good to see real <>BANDS<> on Top Of The Pops,' adds John.*

Without this knowledge, a tagger could, potentially, based on syntactic information alone, select any of nine noun senses.

Likewise, if the tagger didn't know or couldn't guess who The Stones, The Beatles, The Smiths, Hue And Cry or The Who were, they could justifiably assume that they were simply a gang of people: 'a group of people who have a common interest or object, often of a criminal nature':

700231 *WHILE The Stones appealed to the students, and the Beatles to the girls and their mums, The Who were always the lads' <>BAND<>.*

700284 *The Smiths, yeah, they are a thinking man's <>BAND<>.*

700087 *Scots music is all about the voice and the person & dash. That's why Country, folk-rock and older, more emotional forms are so dominant. And, for better or worse, Hue And Cry are a Glasgow <>BAND<>."*

Equally, in the **shake** context below:

700390 *"She believed she was not a lover of women because there was no genital contact." Three weeks before Beauvoir's death, Dr Bair was still <>SHAKING<> her.*

The human taggers' deductive abilities were clear in their choice of disturb:v (to disturb, disconcert or upset the equilibrium of a society, group or person) over move:v (to move someone or something forcefully or quickly up and down or to and fro).

5.3. TAGGERS' WORLD VIEW OR PERSONAL BIAS

How a line was tagged sometimes depended on the tagger's individual 'view of the world'. In the **shake** context below, tags varied depending on whether it was thought that a ghost was a person ("shake off a pursuer") or a thing ("shake off a bad memory").

700176 *A curious combination of two basses, fiddle and accordion meeting the Guardian Women's page. Crawling out from the wreckage of The Cateran, the Joyriders feature two ex-members, Murdo MacLeod and Kai Davidson, plus one tall American named Rick on drums. It takes four songs for them to <>SHAKE<> off their own ghost, but halfway through the aptly named Long Gone it disappears.*

Similarly, at **rabbit**, there were several contexts containing references to 'Roger Rabbit' and 'Peter Rabbit', and tagging varied depending on whether the tagger saw them as toys or animals or neither (in addition, in each case, to them being proper names).³

700090 *Beatrix Potter's Peter <>RABBIT<> is one of Japan's most famous characters: he is often the first Englishman encountered by young readers, and the miniature quality of Potter's stories and illustrations strikes some deep chord in the Japanese heart.*

700240 *The sight of Peter <>RABBIT<> hanging up in an old-fashioned butcher's window brings tears to our eyes, while pretty pink portions prepared and hacked by the supermarket cause no such qualms.*

Similarly, when tagging **onion**, the tagger was faced with a choice between two senses: veg:n "the pungent bulb of a plant . . . , widely used in cooking"; and plant:n "the plant that produces onions". But the matter of when an onion is a vegetable and when it is a plant is a difficult question. For example, when you 'plant the onions', are you putting the bulb (veg:n) in the ground or creating a potential onion plant (plant:n)? And when you 'harvest the onions' or 'lift the onions out of the soil', are they vegetables or still plants?

Since the sense boundaries were blurred, it was necessary to develop a policy and one tagger decided to select plant:n when the onions were still in the soil, had foliage, were being grown, harvested, watered etc., and veg:n when they were being peeled, cooked, sliced etc. However, if I say 'I enjoy growing onions', I surely mean the vegetables not the plants. It seemed that which senses were assigned to the contexts depended on the tagger's personal understanding of when an onion was an onion, and while each tagger developed a policy for their decision-making and could defend their choices, they were keenly aware that another tagger, particularly one who was a keen gardener or cook, could have a different view that was equally defensible.

700095 *Lift the <>ONIONS<> carefully with a fork and lay them out in a sunny place for a few days for them to dry.*

700135 *By August, the foliage will begin to topple and go yellow. Ease a fork under each <>ONION<> to break the roots and leave them on top of the soil to ripen in the sun.*

700028 *Wire mesh, or Netlon stretched from twigs, will also protect the sets from birds and cats. Weed regularly and water thoroughly in dry weather. Your <>ONIONS<> will be ready to harvest in late July or August when the foliage dies and begins to flop over.*

The reportedly personal and largely *ad hoc* nature of taggers' strategies for coping with lexical ambiguity in such cases did not, however, prevent a high level of inter-tagger agreement.

5.4. NON-STANDARD USES OF LANGUAGE

Just as people do not always follow the rules of grammar and syntax, they also use the semantic aspects of language imaginatively and creatively. Beyond the inclusion of recognised figurative sense extensions in the dictionary, there is little provision for this unpredictable aspect of language use.

5.4.1. *Hyperbole, Metaphor and Anthropomorphism*

A problem frequently presented itself when inanimate objects were given human characteristics or emotions.

700254 *Then Olybrius' fury flared and even the ground <>SHOOK<> in fear.*

Only humans or animals can shake with fear and this is made explicit in the dictionary sense: tremble:v “(especially of people or their limbs) to tremble or quiver, especially through fear, illness, shock, or strong emotion”. This deliberate contravention of selectional preferences is used by the author for hyperbolic or humorous effect. This is by no means an uncommon phenomenon in language. While lexicographers attempt to set down generalisations about syntactic or semantic behaviour, identifying constraints and organising linguistic features into manageable categories, language users continue to subvert language for their own ends, be they emphatic, comic, or ironic, or simply because they can.

The human taggers were faced with a choice between tremble:v and move:v “to move (someone or something) forcefully or quickly up and down or to and fro”. The sense move:v would certainly cover *the ground shook*, but since ‘fear’ is the asserted cause of the shaking and is normally restricted to animate objects, it is clear that this is a figurative use and that what is implied at a deeper semantic

level is tremble:v. Should the taggers ignore both what they know to be possible in reality and the semantic features set down in the dictionary entry for tremble:v or ignore the poetic aspect ('in fear') of the context itself and tag it at the literal level? No policy was developed to deal with such cases and the decision was left to the individual taggers. They were also instructed not to confer with each other. The taggers differed in their choices.

A similar case is seen in the use of metaphor in the following **consume** context:

700063 *Apart from the obvious advantage of quieter touring brought by the fifth ratio, the five-speed 'box also seems to have done the SL's fuel consumption no small favour. Overall its exclusively lead-free diet was <>CONSUMED<> at the rate of 10.2mpg, even with a thirsty catalyst as standard.*

Here the author has deliberately taken advantage of the ambiguity between the concrete eat:v "(of a person or animal) to eat, drink or ingest (food, drink or other substances)" sense and the more figurative resource:v "(of people, societies, project, machines, etc) to use up (time, money, fuel or other resources)". An engine is described as 'consuming' a 'diet' of unleaded petrol and having a 'thirsty' catalyst. The language characteristic of the eat:v sense is used to anthropomorphise the engine, but the meaning is the resource:v sense. The human tagger, whilst aware that the context operates on two semantic levels, must choose between the two senses, though neither fully captures what is essentially a concatenation of two senses. Should the tagger assign it a sense according to the language of the imagery or according to the underlying sense? Dictionaries do not allow for metaphor.

This dilemma is echoed in the context below:

700160 *The production will be a flop. In the past couple of years the opposition parties have become skilled at being anti-Thatcherite, <>CONSUMING<> rich pickings from the slow collapse of Thatcherism.*

The imagery is of vultures dining on a carcass, but the actual reference is to political advantages, resources, benefits etc.

A perfect example of an extended metaphor which leaves a human tagger wondering whether to tag the literal use or the actual metaphorical sense, is shown in the context below:

700171 *What was designed by Mrs Thatcher as a Conservative flagship has become, in the words of John Biffen, the Tories' Titanic. Meanwhile, back on the bridge, a new tremor has <>SHAKEN<> the ship with news of a Treasury instruction that low-spending councils must be ready to bail out the high-spenders to reduce the impact.*

On the literal level, a tremor has 'shaken' a ship. But the tremor is a metaphor for the bad news, and the ship is a metaphor for a human institution. Literally, the sense used is move:v, but metaphorically, it is disturb:v.

5.4.2. *Literary References and Idioms*

In the corpus contexts for **bury**, there were three examples of variation on the well-known quotation from Shakespeare's 'Julius Caesar' – 'I come to bury Caesar not to praise him'. In fact, all three instances take the original idiom and capitalize on the ambiguity between the inter:v sense (the original sense intended in the play) and the defeat:v sense.

bury 1 [inter] [vt; often pass] (of a person or group) to deposit (a corpse or other remains of the dead) in a permanent resting place such as a grave, a tomb, the sea, etc., usually with funeral rites.

bury 6.1 [defeat] [vt] to overwhelm (an opponent) totally or beyond hope of recovery.

700107 *Gift's performance will either strike a blow for that much-maligned species, the rock star turned serious actor, or reinforce the opinion that such forays are ego-fueled flights of fancy. No doubt Roland Gift's Shakesperian debut will be attended by critics who will have come not to praise but to <>BURY<> him.*

700132 *It will be her 111th singles match at Wimbledon, one more than Billie Jean King. She has contested 10 finals over the past 18 years. Graf will not be there to praise the American but to <>BURY<> her, just as the 18-year-old Andrea Jaeger annihilated King in their semi-final six years ago, 6-1, 6-1 in 56 minutes.*

As with the metaphorical uses described in the previous section, this use of a popular idiom can be read on two levels, the original or literal sense and the underlying extended sense. The dilemma here would at least give the human tagger cause to hesitate.

5.4.3. *Zeugma*

Another non-standard use of language is seen in the zeugmatic context below:

700028 *Kadar's funeral is the first event to involve workers on a large scale since Mr Grosz replaced him as general-secretary 13 months ago. Mr Pal Kollat, a shipbuilder, described Kadar as an honest man and 'a leader whose lectures we could understand and whose lectures made sense'. The question now is whether the workers respect for the party will be <>BURIED<> along with Kadar.*

The author uses one verb with two nouns, but to each noun a different verb sense applies. While Kadar's burial is literal (inter:v), respect's burial is another, figurative, sense – consign:v (to consign to oblivion or obscurity; to put an end to). It certainly seemed that this context could not be assigned a single sense. This is a further example of the many ways in which language users flout the 'rules' of their language in order to take advantage of its endlessly productive potential.

The various problems encountered by the lexicographers when asked to pair the extremely diverse styles, registers, genres and subject matters covered in a large set of corpus instances with a closed set of dictionary senses are the same problems which humans encounter in their everyday communicative activity. The exercise was carried out under fairly strict time constraints and the lexicographers did not discuss their dilemmas among themselves, neither were they called upon to justify the decisions they made. Discussion of the processes by which such decisions are made is, unfortunately, beyond the scope of this paper.

6. Conclusion

It might be expected, from the extensive catalogue of problematic contexts surveyed in this paper, that the human taggers would have been permanently at odds with each other, and that very little consensus in the sense-tags would have occurred. However, in the total of 8,449 contexts tagged, the rate of agreement was over 95% in most cases (see Kilgarriff and Rosenzweig, this volume).

Almost miraculously, human beings are able to navigate through the multitude of contradictory or mutually incompatible linguistic signals encoded in a text, and with only a small contextual environment as guide, to arrive at a preferred semantic interpretation that is shared by others in their language community. It remains to be seen, from the evaluation of the automatic software tagging results, to what extent the sophisticated techniques employed have managed to approximate to this most human of skills. Can a computer peel an onion?

Notes

* This article is based on a paper given at the SENSEVAL Workshop, Herstmonceux Castle, Sussex, England, 2–4 September 1998.

¹ In addition to the authors, they were Lucy Hollingworth, Guy Jackson, Glennis Pye, and John Williams.

² One of the referees of this paper informed us that these two examples are in fact from a game-theoretic puzzle called 'The Prisoner's Dilemma' for which suggested computational strategies were named 'grudger', 'cheat', 'sucker' etc. 'Tit for Tat' was the strategy that consistently beat all the others!

³ Fully fledged proper names, where there was no relation between any of the word's meanings and its use in the name, were removed from the set of corpus instances to be tagged. However, instances such as 'General Accident' and 'Peter Rabbit', where the word both had one of its usual meanings and was in a name, were tagged with relevant sense and **P** (Proper Name).

References

- Atkins, S. "Tools for Computer-Aided Corpus Lexicography: The Hector Project". *Acta Linguistica Hungarica*, 41(1993), 5–72.
- Atkins, B.T.S. and Levin, B. "Admitting Impediments". In: Zernik, Uri (ed.) *Lexical Acquisition: Exploiting On-Line Resources to Build a Lexicon*, Lawrence Erlbaum, New Jersey. 1991, pp. 233–262.
- Kilgarriff, A. "The Hard Parts of Lexicography". *International Journal of Lexicography*, 11:1 (1998), 51–54.
- Sinclair, J.M., (ed.) *Looking Up: An Account of the COBUILD Project in Lexical Computing*. Collins, London. 1987.
- Stock, P.F. Polysemy. In: Hartmann, R.R.K. (ed.), *LEXeter '83 Proceedings*, Max Niemeyer Verlag, Tübingen. 1984, pp. 131–140.

