Corpus-based Approaches to Logical Metonymy

Term Paper for the Seminar "Preferences in Semantic Processing"

Tatjana Scheffler
tasc@coli.uni-sb.de

April 9, 2003

Contents

1 Introduction

Metonymies yield special problems for semantic processing, as they often lead to type conflicts in the application of predicates and other semantic failures. The important questions posed in the literature include how to detect metonymies, how to determine which of the words or phrases is used metonymically, and what the intended meaning of the metonymy is. If one tries to develop a method of description for metonymies, one further question arises: Is metonymy a conventionalized phenomenon – and should therefore be treated in the lexicon, or is it a productive, context-dependent process – and how should we treat it.

In this paper, we first provide a short introduction into important background literature. In section 2 we recall Pustejovsky's qualia structure as an example for the lexicalization of meaning transfer.

The kernel of this paper is the discussion of two corpus-based analyses that approach some of the above-mentioned questions. Cornelia Verspoor (1997) classified logical metonymies in the BNC according to Pustejovsky's qualia structure. She found that metonymies are largely conventionalized, and can occur only with a small number of words. In section 3 we give a short account of her work.

In contrast, Maria Lapata (2001) tried to extract possible interpretations for logical metonymies from a corpus. Native speaker judgements indicated that metonymy is a very productive process that can only be accounted for by preferences, not by hard constraints on lexical semantics. Her paper is summarized in section 4.

In the last section, a discussion of the proposed explanations follows. It will show that both of them are unsatisfactory, because they do not provide answers for some of the most urgent questions concerning metonymy interpretation, and do not project how these answers can be found. Our particular focus lies on the treatment of metonymy as *preferences*, i.e. defeasible semantic knowledge.

2 Background

2.1 Logical Metonymy

Metonymy is the usage of one word or phrase standing for another, related one (Lakoff and Johnson, 1980). Some examples are:

(1) The *pen* is mightier than the *sword*.

(2) I like to read *Chaucer*. He was a great writer. (Hahn and Markert, 1999)

(3) John began the book. (Verspoor, 1997)

(4) good cook (Lapata, 2001)

(5) good soup (Lapata, 2001)

Example 1 is a classical case of metonymy, where the "pen" stands for the power of the press, and the "sword" for military power. Each of the two abstract concepts is represented metonymically by one characteristic instrument.

Sentences 3-5 are examples of *logical metonymy.* Logical metonymies are the result of polysemous verbs or adjectives such as "begin" or "good". For example, "good" may mean "cooks well" or "tastes good", depending on what noun it modifies. These metonymies arise systematically. The verb "begin" takes an event as a semantic argument. Syntactically, though, it can also combine with a noun. If this noun is not eventive in themself (e.g., such as "operation"), it is still understood as standing for an event. Thus we understand example 3 as being equivalent to "John began reading the book." or maybe "John began writing the book."

The other type of metonymies are *occurrent metonymies*, which can not (or do rarely) occur in different contexts. A typical example of an occurrent metonymy is:

(6) The *ham sandwich* was nice. He gave me a big tip. (M. Egg)

Here, the phrase "ham sandwich" stands for the person that ordered it. It is hardly conceivable of the waitress, if she meets the same person again in the mall two days later, that she says to her friend: "Look, that ham sandwich over there wears a blue tie."

Yet for many metonymical expressions, it is not even clear which of the words in the sentence should be interpreted metonymically. Most authors (Verspoor, 1996; Pustejovsky, 1995) argue that in sentences such as number 3, "the book" is interpreted metonymically, and stands for the event in which it plays an important role. In contrast, Nunberg (1995) speaks in similar cases of "predicate transfer"; i.e. the noun phrase is interpreted literally, but the predicate (e.g. "began") stands for a related predicate (e.g. the one that takes the object of the reading event as an argument).

Figurative language such as metonymy occurs much more often than we usually think. In other words:

> Identity is metonymy of zero distance. (Turner and Fauconnier, 2000)

We will try to argue that the approaches that deal with conventionalized metonymies through the lexicon may be much to rigid. If so, not much is gained by developing mechanisms that can handle these kinds of metonymies. Instead, we need to find out the more general processes that underlie figurative language, and that can be applied to all kinds of metonymy, conventionalized and "occurrent".

2.2 Qualia and Lexical Semantic Structure

Pustejovsky calls *qualia* parts of the semantic structure of lexical entries that are determined by what *function* a concept has in the world, or how a word is *used*. Qualia are for him "aspects of a word's meaning" (Pustejovsky, 1995, p. 76). Four different qualia are specified for nouns, although it is not said whether these are all the roles that can be included in a noun's lexical entry. These qualia are:

- CONSTITUTIVE (the relation between an object and its constituent parts)
- FORMAL (that which distinguishes is within a larger domain)
- TELIC (its purpose and function)
- AGENTIVE (factors involved in its origin or "bringing it about")

Values of a certain quale in the lexical structure of a noun are organized such that words can inherit them from semantic classes they belong to (e.g. all liquids will specify a TELIC role "drink"), but those inherited qualia can in turn be overridden by explicitly specifying another value for a word. (E.g., dictionaries are "consulted" although they are books, which generally are "read".) Not all words have values for all qualia.

As for a word such as "newspaper" many events can be imagined to be specified for certain qualia (because a newspaper is printed matter, as well as a corporation, a physical object, etc.), qualia values are organized in complex type hierarchies, and complex types can be constructed recursively by combining two existing types.

In a metonymic expression, these qualia can be selected to coerce the correct interpretation. If a word specifies more than one event for a certain qualia role, Pustejovsky does not give a decision mechanism as to which one is chosen. Also, not for all words are qualia lexicalized, and Pustejovsky does not provide mechanisms that can deal with these words in any way.

We will discuss the plausibility of qualia in some more detail in section 5.

3 Conventionalized Metonymy

In her paper, Verspoor (1996) argues that logical metonymy is largely conventionalized. She follows Pustejovsky's proposal of introducing qualia in the lexical entries of nouns, which we presented in section 2.2.

3.1 Corpus Analysis

Verspoor looked at the verbs "begin" and "finish" in the Lancaster-Oslo/Bergen (LOB) Corpus and the British National Corpus (BNC), when they combine with a nominal object that is not itself eventive.

(7) John began the book.

(8) John finished the cigarette.

In a sentence such as 7, it is argued that a reading event is missing and must be coerced during semantic interpretation. According to Pustejovsky (1995), this event comes from the qualia structure of the noun "book", where reading is specified as the TELIC role.

In her corpus analysis Verspoor extracted all occurrences of "begin" and "finish" with a non-eventive noun as the object. All the cases were then classified as either choosing the AGENTIVE or the TELIC quale of the noun, or being a contextually dependent metonymy.

3.2 Verspoor's Conclusions

Verspoor found that only very few occurrences of "begin" were metonymic, about 0.41% in the BNC; whereas over 30% of the occurrences of "finish" were found to be non-literal. For both verbs context-dependent metonymies were rare, less than 5% of the metonymic expressions could not be classified as either AGENTIVE or TELIC.

Furthermore, only 20 distinct classes of nouns occurred with their TELIC role selected. Verspoor therefore concluded that logical metonymy is a conventionalized phenomenon, that can only occur with a small number of nouns. She argued that only for those nouns the TELIC quale should be lexicalized, while it is unspecified for other nouns, such as "film" or "nail". That would explain why sentences such as:

(9) John began the film. *(= watching the film)*

(10) John began the nails. *(= hammering them in)*

would be infelicitous, whereas the sentences 7 and 8 work fine. The qualia specified in the lexical entry of a noun would then be available to all kinds of polysemous words, e.g. adjectives like "good" or "fast", and verbs like "begin" or "enjoy". The predicate then selects a certain quale out of the noun's qualia structure in order to produce the correct interpretation of the phrase.

3.3 Discussion

The analysis still lacks important explanations. For a noun like "book", at least two qualia will be lexicalized, because both occur often and are available for metonymic coercions: AGENTIVE and TELIC. How, then, is determined which of these qualia a certain predicate selects in a certain sentence? It is furthermore not the case that a certain predicate always selects the same quale for all its noun arguments. E.g., "begin" would strongly prefer the TELIC quale with the noun "book" (*reading*); but for "house", the TELIC quale is not even possible, and the AGENTIVE must be chosen:

(11) John began the house. *(building the house)*

In general, one could argue that for most artifacts, an AGENTIVE interpretation is possible, because there is commonly a well-defined way they came into being. According to Verspoor's analysis[1], only a few types of nouns will have the TELIC quale specified in their lexical entries. For these, the TELIC interpretation is probably strongly preferred (see example 7), although the AGENTIVE quale is still available for metonymy, if the context supports this. We will argue more for the interpretation of metonymies through preferences in section 5.

In order to render the findings at all useful for the interpretation of metonymies, at least two questions must be answered:

1. When is a quale lexicalized? I.e., what are the conditions, the treshold, etc.?

2. What is the mechanism that selects a certain (part of a) quale in a specific sentence?

4 Interpretation of Metonymies

Lapata (2001) attempted to automatically extract possible interpretations for logical metonymies from corpora, using a statistical method. She concentrated on adjective-noun combinations, such as examples 4 and 5, or:

(12) easy problem

(13) fast horse

She also assumed that an event was missing for the interpretation of those noun phrases, and thus example 12 is basically equivalent to "a problem that is easy to solve"; or example 13 to "a horse that runs fast". In contrast to Verspoor, though, Lapata does not argue that there are only very limited possibilities for the interpretation of these metonymies. Rather, she claims that without context, there is a virtually unlimited number of meanings a metonymic phrase could have. Still, some of these meanings are much more probable than others.

4.1 Model

In order to estimate the probability of a certain interpretation for a metonymy with adjective a and noun n, Lapata tries to find corpus evidence for a sentence where (the adverb corresponding to) a modifies a verb v, and the noun n is either the subject or the object of v. Some adjectives seem to have strong preferences as to whether the metonymic noun is interpreted as the object or the subject of the understood verb. According to Lapata, "easy" triggers a verb-object interpretation, whereas "fast" usually triggers verb-subject interpretations. An adjective like "good" may also allow both verb-object, as well as verb-subject interpretations. A secondary question for Lapata is whether one can extract preference information about the relation (subject or object) an adjective selects.

[1] Her analysis, though, only looked at two different verbs. It is questionable whether her assumption that the same will hold for all logical metonymies is really well-founded.

The probability P for a sentence with verb v, adverb a, and the noun n standing in the relation rel to v can be computed as:

(14) $P(a, n, v, rel) = P(v) \cdot P(n|v) \cdot P(a|n, v) \cdot P(rel|a, n, v)$

Trying to obtain the relevant frequencies from the corpus, one faces a serious sparse data problem. Only one sentence with "fast", "plane" and "fly" could be found in the entire BNC, just as many as for "plane that swoops in fast". Intuitively, "fast plane" should be much more likely to mean "plane that flies fast" than to "swoop in". Therefore, two assumptions were made to estimate the probabilities:

- The likelihood for an adverb to modify a verb is not influenced by this verb's arguments.
- The fact that a certain adverb modifies a verb is independent of the relation rel between this verb and one of its arguments.

Considering these simplifications, the probability P is estimated as follows:

(15)
$$\begin{aligned} P(a, n, v, rel) &\approx P(v) \cdot P(n|v) \cdot P(a|v) \cdot P(rel|n, v) \\ &\approx \frac{f(rel, n, v) \cdot f(a, v)}{f(v) \cdot \sum_i f(v_i)} \end{aligned}$$

4.2 Experiments

Two experiments were then carried out to compare the statistical model with native speaker's judgments. For these experiments, Lapata used a part-of-speech tagged and lemmatized BNC, on which a robust chunk parser was ran in order to extract verb-subject and verb-object pairs.

In the first experiment, Lapata compared her model with 8 examples from the linguistic literature[2]. Table 1 shows the paraphrases of adjective-noun metonymies proposed in the literature along with the most likely interpretations Lapata's model predicted. It can be seen that in many cases ("easy problem", "good shoe", "fast horse", "difficult language", "comfortable chair") the model computed very high probabilities for the intuitive paraphrases. Sometimes it provided synonymic interpretations. Only for the examples "easy planet" and "careful scientist" did the system compute complementary results, although one could argue most of them are also plausible.

Lapata concludes that her model could give a good approximation of the probable interpretations of logical metonymies, but that a randomized test is needed in order to evaluate it further. After all, the examples were chosen for educational reasons in the literature, so they might be a lot clearer and therefore easier than other metonymies of the same kind.

4.3 Human Judgments

Since without context, there is no "correct" interpretation of a given metonymic adjective-noun combination, in the second experiment, the model was compared with human judgments. A good model would rank plausible paraphrases higher than implausible ones, and it was expected that there would occur a linear relationship between human plausibility judgments and the probability ranking the system provides.

Lapata randomly chose nine out of those metonymic adjectives that were listed in (Pustejovsky, 1995) and (Vendler, 1968): *difficult, easy, fast, good, hard, right, safe,*

[2]The examples were taken from (Pustejovsky, 1995) and (Vendler, 1968).

metonymy		interpretation
easy problem	lit.:	solve[a]
	model:	solve, deal with, identify, tackle, handle[b]
easy planet	lit.:	observe
	model:	predict, identify, plunder, see on, work with
good umbrella	lit.:	functions
	model:	cover (SUBJ)
		keep, wave, hold, run for, leave (OBJ)
good shoe	lit.:	wear, walk
	model:	wear, keep, buy, get, stick
fast horse	lit.:	run
	model:	run, learn, go, come, rise
difficult language	lit.:	speak, learn, write, understand
	model:	understand, interpret, learn, use, speak
careful scientist	lit.:	observe, perform, run experiments
	model:	calculate, proceed, investigate, study, analyse
comfortable chair	lit.:	sit on
	model:	sink into, sit on, lounge in, relax in, nestle in

[a] Interpretation given in the literature; see (Lapata, 2001) for reference page numbers.
[b] The five most likely interpretations the model predicted, given in order.

Table 1: Interpretations for examples from the linguistic literature

slow and *wrong*. For each of these adjectives, 10 nouns that occurred in adjective-noun modifications in the corpus were chosen. Then the previously described model was run on the 90 resulting adjective-noun pairs. To achieve this, the corpus was parsed with Gsearch (Corley et al., 2001) to retrieve verb-subject and verb-object pairs. The model yielded a probability ranking of all possible verbs, which (for each adjective-noun pair) was split in three "probability bands" of equal size: High, Medium and Low. This was done twice, once for each type of relation (subject or object).

For each probability band, one interpretation (= verb) was randomly selected, so that a complete number of 90 (adj-noun pairs) $\times$ 2 (relations) $\times$ 3 (probability bands) = 540 stimuli resulted. As these were too many to be judged by native speakers, for each adjective, only 5 of the 10 nouns were selected for subject and object interpretation, respectively. The final 270 stimuli were presented to two different groups of subjects, 135 stimuli each. The subjects were asked to judge the appropriateness of a certain paraphrase to a given adjective-noun pair.

Results

The comparison was evaluated using variance and correlation analyses techniques. The variance analysis revealed that the differences in the appropriateness rating across the probability band was significant. In particular, examples from the "High" probability slice were rated significantly better than examples from the "Medium" or "Low" slices. This effect was further verified by the correlation analysis, which proved that there was a significant correlation between the individual probabilities of interpretations the model predicted and the interpretation preferences by human subjects. The correlation between the model and the human judgments was $r = .53$ for verb-object interpretations and $r = .21$ for verb-subject interpretations. Intersubject agreement, which can be seen as the upper bound for system performance in this case, was .67 and .65 for the two groups, respectively.

As can be seen, the model performed quite well on verb-object paraphrases, especially if one takes into account that human subjects did not agree very well in their judgments, either. The correlation for verb-subject paraphrases is much worse, which Lapata attributes to the fact that the parser performs much better in the extraction of verb-object pairs, and that most of the verbs in the set had verb-object biases.

Naive Baseline

Finally, Lapata compared her results with those of a naive baseline that did not take the dependency between adjective (adverb) and verb into account. This would, apparently, be Verspoor's approach, as she argued that the lexicalized qualia are available for all kinds of different metonymic predicates in the same way. Following this analysis, one should predict that the interpretation of logical metonymies is independent of the polysemous predicate (here, the adjective) that triggers it.

The naive model thus ignores the adjective and just takes the relation between the noun and the verb into account. The probability of an interpretation is now computed as:

(16) $$P(v|n, rel) = \frac{f(v, n, rel)}{f(n, rel)}$$

Whereas the former model (that included the adjective) yielded a Pearson correlation coefficient of $r = .40$, the naive model achieves much lower correlation, only $r = .25$ (.29 for verb-object interpretations; no correlation could be found for verb-subject paraphrases). This difference to the more complicated model is statistically significant.

4.4 Conclusion

With her experiments, Lapata showed the possibility to extract interpretations for adjective-noun metonymies from a corpus. Furthermore, the possible interpretations for a metonymy can be ranked automatically by probability, and this ranking was shown to correspond largely to human *preferences*.

Lapata also demonstrated that the common claim promoting qualia, that certain events are lexically associated with certain nouns, and that those events are available across metonymic expressions, can not be true in general. The naive baseline that followed such an approach performed significantly worse than Lapata's model that took the individual adjectives (the elements that evoked the metonymy) into account.

5 Discussion: Metonymy and Semantic Preferences

In this paper we presented two attempts to make use of corpus data in the approach to the interpretation of logical metonymies. Logical metonymy was characterized as *conventional* or *regular* metonymy, arising from the inherent polysemy of certain adjectives ("good") and verbs ("begin").

While Verspoor argued with Pustejovsky that metonymy interpretation is very restricted, Lapata sees almost unlimited possibilities. Both found evidence for their claim, from a corpus or human plausibility judgments. This disagreement points to one dimension of the metonymy interpretation task: to what extent are metonymies lexicalized, and what role does context play.

5.1 Lexicalization

Lexicalization of possible metonymic interpretations as "qualia" seems not to cover many cases. Kilgarriff (2001) carried out a small study on a number of randomly chosen words to see whether their metonymic usages could in principle be explained with appropriate qualia structures. He found that for most examples, this was not the case. Kilgarriff concluded that not only the context, but also further world knowledge, and especially literary and cultural history (what has been said before) play an important role in the interpretation of metonymy.

Of course, Pustejovsky as well as Verspoor both mention that qualia structure is only supposed to be applied to most, conventional cases. The context can always override these lexicalized constraints and trigger a different interpretation, such as for example in 17 (b). However, it is not evaluated how this interpretation can be found, and how the failure of a conventional metonymic interpretation can be detected.

(17) (a) My goat eats anything.

(b) After he finished the grass, he *began the book*.

A further problem arises if more than one interpretation is lexicalized. For example, it is not known whether the TELIC or AGENTIVE quale of "book" would be preferred for a sentence as 7 (meaning "John began to read" or "John began to write", respectively), and whether this preference is generally valid.

In conclusion, the proposed qualia structure and its use in the interpretation process for metonymies seem to be much too restricted to account for an interesting number of cases. Furthermore, the many unanswered questions for the application of the theory render it unusable for any practical application that produces interpretations for such sentences. It seems indispensable to reconsider metonymic interpretations in a *preferential* algorithm, i.e. in a way that leaves the possibilities very open for context specification, as well as yields a preferred analysis for more conventional cases.

5.2 Preferences and Probability

For this purpose, Lapata's approach that uses statistical methods to extract possible paraphrases for metonymy from corpora is a good first step. It will provide us with a ranked list of more and less probable interpretations for specific metonymic expressions. This list is certainly needed for an analysis that makes use of semantic preferences. Still, her algorithm does not take the context into account. Although in principle the computed plausibility hierarchy may be valid (and verified by human judgments), in a specific context the ranking may be very different, and some of the first (most probable) interpretations may not be possible at all. A mechanism that choses between the interpretations in the top-most slice of the list is the next necessary step. This mechanism needs to take further context information into account, such as for example selectional restrictions a coerced predicate may have.[3]

Lapata's approach does also not make strong assumptions as to which element in the metonymic interpretation is non-literal and needs to be reinterpreted. In our view, neither of the two proposals (the one that the noun receives metonymic interpretation through its qualia structure; and the other that the nouns are literally interpreted, and the predicate is transferred to a related one (Nunberg, 1995)) has completely convincing arguments in its support. Predicate transfer seems to be a more unrestricted process that is also better suited to handle some anaphoric

[3]In the "goat" example, the conventionalized "read" requires its subject to be capable of reading, which a goat is not. Therefore, this interpretation is impossible.

relations cooccurring with metonymic expressions. An approach such as Lapata's, that seeks the (re-)interpretation in the combination of both predicate and noun, is a good choice at this point.

5.3 Metonymic vs. Literal Interpretation

One dimension of the problem that none of the quoted authors addressed is the distinction between cases of metonymic and literal interpretation. Not only do different metonymic interpretations of the same phrase compete with each other, but they also compete with a literal interpretation. It is not always the case that a non-literal meaning is accompanied by a semantic type clash.

(18) I like Chaucer.

For example for verbs with very soft selectional restrictions such as "like", no type clash occurs in example 18. In this case, one could argue that the reinterpretation is triggered because of world knowledge, as Chaucer is dead and it is unlikely that someone still knows the person. Even so, one could like a dead person. But if I say:

(19) I like Günther Grass.

it is also much more probable that I mean his works than the person. Here, a metonymic interpretation is preferred over a literal one, although no semantic clash would lead to reinterpretation in an algorithm that always first tries to interprete sentences literally. Markert and Hahn (1997) dealt with this problem in connection with anaphors. But no systematic analysis of the factors that lead to the preference of different metonymic or literal interpretations over one another are known to us. At this point, full world knowledge can not be expected to be available for algorithms that interpret sentences. Therefore, the most important aspects and those that cover the most phenomena must be discovered and made usable.

5.4 Corpus Analyses and Further Work

Corpus analyses can help to find new approaches to preferences in the interpretation of (possibly) metonymic expressions. Extraction of possible interpretations for metonymies from corpora using approaches such as Lapata's (or Kilgarriff's) will determine the range that analyses have to cover.

Then, mechanisms have to be developed that determine under which circumstances one metonymic interpretation is preferred over the other. Possible indicators are:

- anaphoric reference
- semantic type clashes
- collocational information from other words in the sentence (e.g. the subject)
- world knowledge

Finally, it has to be determined under which circumstances a metonymic interpretation is preferred over a literal one. The types of information that are useful for this purpose are in principle the same as above. It must also be established whether there is a general preference for literal interpretation, or whether this can not be assumed. It would be possible to conduct a corpus analysis to find out how literal and metonymic interpretations statistically relate.

5.5 Conclusion

In this paper, we presented two approaches to metonymy analysis that used corpora as evidence. The approaches differed in their respective assumptions that logical metonymies are largely lexicalized, or are context-dependent and can be captured by statistical methods. We argued that a strict lexicalization of interpretation possibilities is much too restrictive, although some metonymies are more conventional than others.

We sketched some questions that have been addressed in the literature only in part, and how the range of possible answers looks like. Finally, we projected how one might go about answering some of these questions.

References

Ann Copestake and Ted Briscoe. Semi-productive polysemy and sense extension. *Journal of Semantics*, 12(1):15–67, 1995.

Steffan Corley, Martin Corley, Frank Keller, Matthew Crocker, and Shari Trewin. Finding syntactic structure in unparsed corpora: The gsearch corpus query system. *Computers and the Humanities*, 35:81–94, 2001.

Udo Hahn and Katja Markert. On the formal distinction between literal and figurative language. In *Portuguese Conference on Artificial Intelligence*, pages 133–147, 1999.

Adam Kilgarriff. Generative lexicon meets corpus data: the case of non-standard word uses. Technical report, ITRI, University of Brighton, 2001.

George Lakoff and Mark Johnson. *Metaphors We Live By.* University of Chicago Press, 1980.

Maria Lapata. A corpus-based account of regular polysemy: The case of context-sensitive adjectives. In *Proceedings of NAACL 2*, 2001.

Katja Markert and Udo Hahn. On the interaction of metonymies and anaphora. In *Proceedings of IJCAI*, pages 1010–1015, 1997.

Geoffrey Nunberg. Transfers of meaning. *Journal of Semantics*, 12:109–132, 1995.

James Pustejovsky. *The Generative Lexicon.* MIT Press, Cambridge, MA, 1995.

Mark Turner and Gilles Fauconnier. Metaphor, metonymy and binding. In Antonio Barcelona, editor, *Metaphor and Metonymy at the Crossroads*, volume 30 of *Topics in English Linguistics.* Mouton de Gruyter, 2000.

Zeno Vendler. *Adjectives and Nominalization.* Mouton, The Hague, NL, 1968.

Cornelia M. Verspoor. Lexical limits on the influence of context. In *Proceedings of the 18th Annual Conference of the Cognitive Science Society*, San Diego, CA, 1996.

Cornelia M. Verspoor. Conventionality-governed logical metonymy. In *Proceedings of IWCS 2*, Tilburg, NL, 1997.