


In Support of the Equal Rights Movement for Literal and Figurative Language – A Parallel Search and Preferential Choice Model

Udo Hahn & Katja Markert

 Computational Linguistics Research Group
Freiburg University, Werthmannplatz 1
D-79085 Freiburg, Germany
<http://www.coling.uni-freiburg.de/>

Abstract

We challenge the commonly held view that the interpretation of metonymies should proceed from a literal-meaning-first approach and argue for an equally balanced treatment of literal and figurative language use. Resulting ambiguities are handled by a combination of two techniques. First, we incorporate discourse constraints into metonymy resolution, reflecting the systematic interaction patterns between the resolution of nominal anaphora and metonymies. We, second, impose constraints on metonymies that are based on pragmatic criteria.

Introduction

One of the hot natural language topics discussed in the field of cognitive science is the skewed relation among literal and figurative language such as metaphors, metonymies or irony (for a survey, cf. Gibbs (1994)). The traditional literal-meaning-first hypothesis (henceforth, LMF), which still prevails in the field of (computational) linguistics and philosophy, subscribes to the view that the figurative meaning of an utterance is derived from its literal one. The ensuing approach to understanding figurative language is a two-stage *serial* model where the figurative meaning is computed only after the literal meaning of an utterance has been processed and considered a violation of communicative norms (cf. Grice (1989)).

There is, however, a growing body of experimental data from cognitive studies (Gibbs, 1989; Gerrig, 1989) that cast doubt on the legitimacy of LMF. In addition, Recanati (1995) has made a strong attack against LMF with convincing though not experimentally backed up arguments concerning its philosophical underpinnings. Our contribution to this discussion consists of a computational model for metonymy resolution. This model has grown out of our experiments in automatic text understanding (Strube & Hahn, 1996; Markert et al., 1996) and — after careful examination of the relations holding in written (expository) texts among anaphora and metonymies — is grounded on three major claims:

- Literal and figurative interpretations are computed *in parallel*. Specifically, the computation of a figurative interpretation of an utterance is not dependent on the previous processing of a literal one. Hence, LMF is rejected.
- *Discourse context* is a crucial determinant for selecting among competing literal and figurative interpretations (cf.

Gibbs (1993) and Keysar (1994) for experimental support of this claim). We further supplement these discourse constraints by pragmatic criteria for metonymy interpretation.

- Unless expensive backtracking is granted, common LMF approaches yield entirely inadequate results under *incremental* parsing conditions, i.e., perform counterintuitive when applied to the usual mode of human parsing behavior.

The Interaction of Metonymies and Anaphora

An almost canonical definition considers an expression *A* a *metonymy*, if *A* deviates from its literal denotation in that it stands for an entity *B* which is not expressed explicitly but is conceptually related to *A* via a metonymic relation *r*. Most studies on the processing of metonymies subscribe to an LMF approach where metonymy resolution is triggered only when a sortal conflict (or, deriving from that, a syntactic anomaly) is detected in the course of the interpretation of an utterance.¹ Such an approach disregards the textual embedding of most metonymies, which is reflected by the systematic interaction patterns in the resolution of nominal anaphora and definite metonymic noun phrases. This is surprising, since metonymy is generally regarded as a phenomenon of reference. The subsequent discussion is intended to reveal various interaction patterns for metonymies and anaphora and the impact they have on the validity of LMF. For a definite noun phrase *A*, we may distinguish between the following usage patterns:
I. A is anaphoric and ...

1. *A* is *not metonymic*.² Hence, the literal denotation of *A* allows for anaphora resolution; in addition, *A* fulfils the required sortal constraints.
2. *A* corresponds to a *predicative metonymy*.³
 - (i) Ten minutes before the *notebook* switches off, it starts beeping.
 - (ii) The clock frequency of the *computer* is reduced to 8 MHz.

¹In the majority of LMF models, the violation of communicative norms is narrowed down to the violation of *intrasentential* well-formedness conditions. It is the latter understanding of LMF we refer to in the remainder of this paper.

²We restrict figurative speech in this paper to metonymies.

³Stallard (1993) and Nunberg (1995) discern cases of predicative and referential metonymy depending on whether the literal or the intended referent is available for subsequent pronominal reference.

In (ii), a literal denotation is available for anaphora resolution in (i) - “computer” resolves to “notebook”. It does not, however, fulfil all sortal constraints, as “clock frequency” is a property attributed to a processor, not to a computer as such. No significant interaction between anaphora and metonymy resolution need take place.

3. A corresponds to a *referential metonymy*.

- (i) We also tested *the printer Epson EPL-5600*.⁴
- (ii)a) I liked *the laser*, as its printouts were excellent.
- b) I liked *the laser*.

In both cases, anaphora resolution is fully dependent on metonymy resolution. For case a), the resolution of the part-for-whole metonymy, “laser” for “laser printer”, can be achieved without information about the possible anaphoric antecedents of “the laser”, as the analysis of the sentence reveals a sortal conflict (the combination with “its printouts” fails). Thus, in a syntax-first approach, after syntactic processing a (quite sophisticated) metonymy resolution procedure had to precede anaphora resolution to solve this case. Considering an incremental approach, anaphora resolution for “the laser” would be triggered *before* the information about the sortal conflict were available so that the problem is reduced to case b). Example b) constitutes a counter-example to LMF approaches, as the literal meaning of the sentence is not semantically or syntactically deviant but a metonymic interpretation is obviously intended. Only the information about possible anaphoric antecedents of “the laser” can help finding this metonymic reading at all. In addition, this information may help with choosing amongst several metonymic readings — excluded are those readings, which do not allow for anaphora resolution (e.g., the competing metonymic reading “laser” for “light”).

II. A is not anaphoric and ...

1. A is not metonymic. When the literal denotation of A fulfils all sortal restrictions but does not allow for anaphora resolution the necessary criteria for *functional anaphora* (Markert et al., 1996) are fulfilled:

- (i) We also tested *the printer Epson EPL-5600*.
- (ii) I did not like *the paper-tray*.

The parallel structure of example 1.3.(ii).b shows that these criteria can also be met by anaphoric noun phrases A which are referential metonymies. In order to distinguish between these two cases (functional anaphora vs. referential metonymy) we will incorporate additional pragmatic constraints for metonymy resolution later on.

2. A is metonymic. In this case, predicative and referential metonymies are either marked by sortal conflicts or can only be recognized by very sophisticated inferential mechanisms which are not an issue here.

⁴Given the original discourse context of this text fragment, “Epson EPL-5600” is already known to be a laser printer.

We argue for an integrated model that accounts for the systematic *interdependencies* between nominal anaphora and metonymies. Metonymy resolution is not only crucial for anaphora resolution. *Discourse* restrictions may even facilitate metonymy resolution and the resolution of corresponding ambiguities. In addition, evidence from preceding discourse elements may override LMF by referential constraints. These interaction effects are highly rewarded for (though by no means restricted to) incremental approaches to natural language analysis, as sortal conflicts may not be recognized at the time when anaphora resolution is carried out.

Basic Conceptual and Semantic Constraints

The parser we use establishes syntactic structures only if conceptual and semantic constraints between the lexical items involved are met. Conceptual checks identify well-formed role chains between the concepts denoted by the lexical items; semantic checks determine whether these chains mirror metonymic relationships or literal ones. The representation structures to which these checks refer are grounded in a hybrid terminological knowledge representation framework (cf. Woods & Schmolze (1992) for a survey). The concept hierarchy consists of a set of concept names $\mathcal{F} = \{\text{COMPUTER-SYSTEM, PRINTER, ...}\}$ and a subclass relation $isa_{\mathcal{F}} = \{(\text{LASER-PRINTER, PRINTER}), (\text{NOTEBOOK, COMPUTER-SYSTEM}), ...\} \subset \mathcal{F} \times \mathcal{F}$. The set of relation names $\mathcal{R} = \{\text{has-physical-part, has-laser, clock-frequency-of, ...}\}$ contains the labels of all conceptual roles. These are also organized into a hierarchy by the relation $isa_{\mathcal{R}} = \{(\text{has-laser, has-physical-part}), (\text{clock-frequency-of, property-of}), ...\} \subset \mathcal{R} \times \mathcal{R}$. For every conceptual instance i its (direct conceptual) class is given by $class(i) = C$ or by $i inst-of C$.

Conceptual Relatedness

By definition, we associate with a lexical item *lex* a standard denotation in terms of a *concept*, $LEX.C \in \mathcal{F}$ (e.g., the lexical item “computer” denotes COMPUTER-SYSTEM). Furthermore, $lex.r$ refers to the *instance* corresponding to *lex*.⁵ If a syntactic link between two lexical items, A and B, is to be allowed, the concepts A.C and B.C must be conceptually related. In order to determine conceptual relatedness, we employ a *path finder*, which performs an extensive search in the domain knowledge base looking for *well-formed paths* between two concepts. We will not go into the formal details of well-formedness criteria for a conceptual path $(r_1 \dots r_n)$ ($r_i \in \mathcal{R}$) linking two concepts $x, y \in \mathcal{F}$ (cf. Markert & Hahn (1997)). Rather, we only briefly mention that complete connectivity (compatibility of domains and ranges of the included relations) and non-cyclicity (exclusions of inverses of relations) are required for a *conceptually well-formed path*.⁶

⁵If *lex* is meant literally, $lex.r \in LEX.C$ holds, whereas for metonymies this need not be the case.

⁶The non-cyclicity condition already limits the search space of the path finder considerably. We are currently going to extend the path-finding algorithm by additionally incorporating discourse constraints from the very beginning of the search.

Literalness vs. Figurativeness

Every well-formed path between A.C and B.C is interpreted by the *path evaluator*. Certain predefined path patterns are used to distinguish between a subset \mathcal{L} of all types of well-formed paths, which is labeled “*literal*”, another subset \mathcal{M} labeled “*metonymic*”, and all remaining paths. Hence, a literal path between A.C and B.C mirrors a literal interpretation of both A and B, whereas a metonymic path between A.C and B.C mirrors a metonymic interpretation of A or B.

Literal Paths. We call a relation chain a *literal* one, if it can be treated as a single relation. All paths of unit length I are included in \mathcal{L} , as they are “literal”, by definition (they refer to the conceptual roles directly associated with a concept definition). In addition, we incorporate empirical observations about the transitivity of relations. Chaffin (1992) distinguishes several subtypes of *part-whole* relations and claims that any of these *subrelations* are transitive, while the general *part-whole* relation usually is not. Thus, a relation chain containing only relations of one of these subtypes is again *a relation of the same subtype*, whereas a relation chain containing several different types of *part-whole* relations does not constitute a *part-whole* relation any more. Accordingly, we have included the path patterns (*has-physical-part**), (*collection-member**), (*mass-portion**), (*process-phase**), (*event-feature**), (*area-place**) and the corresponding inverses like (*physical-part-of**) in \mathcal{L} .

Metonymic Paths. Following established classifications (Lakoff, 1987; Fass, 1991) we incorporate *whole-for-part*, *part-for-whole*, *producer-for-product*, *container-for-contents* and *material-for-object* metonymies. We will not go into any technical details how the path patterns for these metonymies can be determined (cf. Markert et al. (1996)), and consider instead example I.2.(ii), where a syntactic link between “*clock frequency*” and “*computer*” is checked by searching for a well-formed path between the corresponding concepts CLOCK-FREQUENCY and COMPUTER-SYSTEM. The path finder locates a single, metonymic path (*clock-frequency-of-cpu-of-motherboard-of*). We may deduce by this path pattern that “*computer*” denotes a *whole-for-part* metonymy for an instance of the concept class CPU (*cpu-of* as well as *motherboard-of* denote metonymic “*part-of*”-like relations which allow “*computer*” to stand for CPU, though the latter is not expressed explicitly). Thus, although we take the literal denotation of the *lexical item* “*computer*” as the starting point for our analysis, the computation of the metonymic interpretation of the *utterance* is in no way dependent on computing a deviant literal interpretation of the *utterance* first.

Assuming a strict preference for literal over metonymic path patterns (as characteristic of LMF approaches) does not penalize metonymy resolution in this example, since only a metonymic path between CLOCK-FREQUENCY and COMPUTER-SYSTEM is found. In contradistinction, example I.3(ii).b demonstrates that this ranking can be overridden by contextual information in the case of definite noun phrases. Furthermore, in an incremental approach a sortal conflict may

be detected at a rather late stage of processing (cf. example I.3(ii).a). Hence, expensive backtracking becomes necessary, if information about possible anaphoric antecedents is not taken into account early on. As a consequence, we prefer literal over metonymic paths only in those cases when both concepts are expressed by *indefinite* noun phrases, verbs or adjectives. If a *definite* noun phrase occurs, we will incorporate information about possible anaphoric antecedents.

Discourse Constraints

When a definite noun phrase A in the utterance U_i can syntactically be bound to an H with a conceptual correlate H.C (e.g., the main verb of the sentence), anaphora resolution for A is triggered. Based on a functional centering framework (Strube & Hahn, 1996) each utterance U_{i-1} is assigned a set of forward-looking centers $C_f(U_{i-1})$ containing the possible anaphoric antecedents for an anaphor in U_i . Thus, $C_f(U_{i-1}) = [c_1.r, \dots, c_n.r]$ must be searched for potential referents of the expression A in U_i , allowing for coercions of A .

Anaphora and metonymy resolution proceed as follows.⁷ First, the path finder and evaluator are called with H.C and A.C as their arguments. This returns a list of all well-formed paths between H.C and A.C, which are marked for their literalness or figurativeness. As shown above these paths are computed in parallel and independently of each other. In the example I.3.(ii).b, anaphora resolution for “*the laser*” is triggered when the syntactic link between “*laser*” and “*liked*” is to be established. The path finder and evaluator are called with the corresponding concepts, LASER and LIKE, returning the following two paths:

1. (*like-patient*)⁸ is a *literal* path expressing that “*laser*” is meant literally.
2. (*like-patient has-laser*) is a metonymic path expressing a *part-for-whole* metonymy where “*laser*” stands for an instance of the class LASER-PRINTER =: B.

We then determine how these findings combine with anaphora resolution, i.e., whether a specific interpretation leads to a referent $c.r \in C_f(U_{i-1})$ for A , thus integrating discourse constraints into the selection among competing figurative and literal interpretations.

When A is *literal* and *anaphoric*, a literal path between A.C and H.C exists and an element is in the centering list whose class is a subconcept of A.C.

When A is a *metonymy* for an instance of the concept B via the relation r , there exists a metonymic path containing r between A.C and H.C. When A is *anaphoric* we have to distinguish between the cases of a predicative metonymy and a referential one. Whereas for a *predicative metonymy* A.C is available for anaphora resolution, for a *referential metonymy* only B is.

⁷ We here give a short outline of our algorithm only. For technical details and a more formal description see Markert & Hahn (1997).

⁸ The ACTION part of the concept hierarchy carries case-role-style specifications. For instance, *like-patient* characterizes the patient of a LIKE activity (in this case, the *laser (printer)*).

With regard to our example *I.3.(ii).b*, the literal interpretation of *A* and the one considering it a predicative metonymy do not allow for anaphora resolution, as $C_f(U_{i-1}) = [Epson\ EPL-5600]$ with $class(Epson\ EPL-5600) = LASER-PRINTER$ holds. However, if we regard *A* as a referential metonymy, anaphora resolution succeeds, as $class(Epson\ EPL-5600) = LASER-PRINTER\ isa_{\mathcal{F}}^* B = LASER-PRINTER$ holds. Choosing among these readings, we propose a ranking with respect to the ordering relation *str* as in Table 1.

Table 1: Ranking Constraints

literal(A) \wedge anaphoric(A)
$>_{str}$ predicative metonymy(A) \wedge anaphoric(A)
$>_{str}$ (referential metonymy(A) \wedge anaphoric(A))
\vee (literal(A) \wedge \neg anaphoric(A))
$>_{str}$ metonymic(A) \wedge \neg anaphoric(A)

Considering our example, we encounter a case of ambiguity as indicated in line 3 of Table 1. This results in instantiating two separate readings of sentence *I.3.(ii).b* and discarding the *predicative* metonymy reading corresponding to the last line in Table 1.

Metonymies marked by a sortal conflict are handled as in LMF approaches, since no competing literal interpretation exists. Thus, “computer” in example *I.2.(ii)* is treated as a *whole-for-part* metonymy.⁹ As the literal meaning of “computer” allows for anaphora resolution, we treat the metonymy as a predicative one (second line of Table 1). In contradistinction to LMF approaches, our algorithm also allows for the resolution of anaphoric metonymies that are not marked by a sortal conflict as example *I.3.(ii).b* shows. In addition, we do not depend on information about sortal conflicts that may not be available at the time point when metonymy or anaphora resolution is called for. As a consequence, an incremental approach to metonymy resolution is becoming more feasible.

Similarly, anaphora resolution proceeds as usual when a literal interpretation allows for anaphora resolution (first line of Table 1). Unlike LMF, nominal anaphora that are referential metonymies can be resolved incrementally without the need for backtracking. So, in example *I.3.(ii).b* “the laser” can be resolved to “Epson EPL-5600” in sentence (i) in at least one reading (third line of Table 1). Thus, different referential mechanisms are at work in the interpretation of definite noun phrases — namely, anaphoric reference by using a more general expression (*printer* for *laser printer*) and anaphoric reference by using a metonymic expression (*laser* for *laser printer*) — and they are covered by the same algorithm.

Additional Pragmatic Constraints

A crucial problem with our approach lies in the disjunction in line 3 in Table 1, which is a continuous source of ambiguity. Whereas it leads correctly to an ambiguity in example *I.3.(ii).b* as explained above, it also leads to the same kind of ambiguity in example *II.1.(ii)* where only the literal

⁹Note that the *processing* of these metonymies is nevertheless different from LMF approaches, as we do not compute a deviant literal interpretation of the utterance, first.

functional anaphor reading is intended. The interpretation of “paper tray” as an anaphoric reference to “the printer Epson EPL-5600” via a referential metonymy should be discarded. The algorithm as proposed so far does not yield any criteria to prevent this kind of overgeneration. It is important to note that this is neither an artificial problem created by our algorithm in the first place nor one dependent on accepting or rejecting LMF. Instead, it mirrors the fact that the resolution of metonymies is not fully constrained by discourse and metonymic patterns like *part-for-whole* and that some metonymies are more conventional than others. The question which arises is why the relation *has-laser* between LASER-PRINTER and LASER can be easily exploited for a referential metonymy, whereas given the relation *has-paper-tray* between LASER-PRINTER and PAPER-TRAY, this is not that likely. We investigated three heuristics to further constrain coercion by pragmatic constraints.¹⁰

Singular relationships. First, we argue that a relation *r* linking the concept B to A is more likely to be used for metonymic purposes, if it has the number restriction *exactly one*, i.e., if *r* is a function from B to A. In the example *I.3.(ii).b*, “laser” is more likely to be a metonymy for “laser printer” than “paper tray” (*II.1.(ii)*), as a laser printer is normally supplied with exactly one laser, whereas it may have several paper trays. Thus, *singularity* often marks the importance of certain relationships. *Producer-for-product* metonymies exploiting the relationship *produced-by*, e.g., fulfil this criterion by definition, whereas *part-for-whole* relationships do not. This might possibly explain why *producer-for-product* metonymies are much more conventionalized.

Typical relationships. Second, a relation *r* between B and A is more easily used for metonymic relationships, if it characterizes an instance of B as a typical member of its class. A member of the class LASER-PRINTER receives its typicality by having a LASER in contrast to other printers. Therefore, the relation “*has-laser*” can be easily exploited for metonymic references. We approximate the notion of typicality by characterizing a relationship between B and A as a “*typical*” one whenever it sets B apart from at least one of its direct superconcepts.

Familiar relationships. Our third heuristic concentrates on the relationships between *instances* in the knowledge base already known from previous processing or from *a priori* knowledge. Let us assume that an expression A is a referential metonymy, i.e., $class(A.r) = B$ where $B\ isa_{\mathcal{F}}^* A.C$ does not hold. B and A.C are related via a metonymic relation *r*. We found that this is more likely, if A.r (*inst-of*B) already has a filler belonging to A.C for the role *r*. So, if we talk about a laser printer that already has a filler for its role “*has-laser*”, this role is more likely to be used metonymically later on as the reader is already *familiar* with it. This criterion is almost always fulfilled for *producer-for-product* metonymies, as the producer of a specific product is normally known either from

¹⁰We here concentrate on referential metonymy, as predicative metonymy is normally indicated by a sortal conflict.

the name of the product or from previous knowledge.

Summing up, the ambiguity in line 3 of Table 1 is only established when the metonymic interpretation fulfils at least two of the three criteria above. Otherwise, only the literal interpretation is created. Hence, for example *I.3.(ii).b*, both interpretations are created (as the metonymic relation “*has-laser*” fulfils criterion 1 and 2), whereas in example *II.1.(ii)* only the literal interpretation is, as the supposed metonymic relation “*has-paper-tray*” fulfils none of the above criteria.

Evaluation

The test set for our evaluation experiment was composed of naturally occurring texts, *viz.* 26 German product reviews from the information technology domain. The main part of the evaluation was carried out manually in order to circumvent error chaining in the anaphora resolution, while path finding and evaluation was done automatically. The sample contained 103 metonymies, 291 nominal anaphors, 351 functional anaphors, all occurring in 606 utterances. Table 2 contains the quantitative distribution of occurrences of metonymies in our test set. The columns indicate whether a sortal conflict (*s.c.*) occurred directly (known at the time of anaphora resolution and, hence, resolvable in an incremental framework), indirectly (not known at that time), or not at all. With 42.7%, direct sortal conflicts are below expectations, but LMF approaches face even more serious problems for 29.1% of the metonymies which are not marked by any sortal conflict at all. We also want to point out the significant rate of co-occurrences of metonymies and anaphors (56.3%). Thus, anaphoric processes indeed seem to facilitate the occurrence of metonymies, an observation that supports our claim to account for the systematic interaction of both processes.

Table 2: Distribution of Metonymic Noun Phrases

	direct s.c.	indirect s.c.	no s.c.	Σ
Anaphors	15 (14.6%)	27 (26.2%)	16 (15.5%)	58 (56.3%)
Other def. NPs	7 (6.8%)	2 (1.9%)	5 (4.8%)	14 (13.6%)
Indef. NPs	22 (21.4%)	0 (0%)	9 (8.7%)	31 (30.1%)
Σ	44 (42.7%)	29 (28.2%)	30 (29.1%)	103 (100%)

Table 3: Resolution Rates for Metonymic Noun Phrases

	direct s.c.	indirect s.c.	no s.c.	Overall Rate
Anaphors	93.3%	85.2%	87.5%	87.9%
Other def. NPs	100.0%	0.0%	0.0%	50.0%
Indef. NPs	72.7%	0%	0%	51.6%
Overall Rate	84.1%	79.3%	46.7%	72.8%

Table 3 depicts the resolution rates for our approach.¹¹ The high proportion of correctly resolved anaphoric metonymies (87.9%) lends some credit to our ranking in Table 1. Especially important is the fact that 85.2% of all anaphoric metonymies marked by an indirect sortal conflict and 87.5%

¹¹We cannot compare our results to those from other approaches, since none of them has been empirically tested, so far. We, therefore, incorporate general observations of the problems LMF approaches handling only metonymies with sortal conflicts are likely to face. Syntactic irregularities played a decidedly minor role in our test set and are therefore excluded from the evaluation.

without a sortal conflict can be resolved by an incremental approach without backtracking – due to the timely consideration of anaphoric constraints. LMF approaches are systematically limited to account for only 42 anaphoric metonymies (72.4%) marked by a sortal conflict (cf. Table 2, anaphors marked by a direct (15) and indirect (27) sortal conflict, respectively).

Our resolution rates are considerably worse for definite non-anaphoric noun phrases (50.0%) and indefinite noun phrases (51.6%), the main reason being that metonymies without a sortal conflict are not recognized, as anaphoric constraints are not applicable. Hence, LMF approaches are insufficient for non-anaphoric noun phrases as well. However, we do not have a solution to this problem either

As shown in Table 2, 58 from 291 anaphors (19.9%) are metonymic. As 47 of them are referential metonymies, a 16.2% increase of anaphora resolution capacity can be achieved through the incorporation of metonymy resolution.¹² None of the LMF approaches can resolve the 16 metonymic anaphors (constituting 5.5% of all anaphors), which are not marked by a sortal conflict.

Considering the impact of overgeneration, 351 cases of functional anaphora have to be considered (e.g., example *II.1.(ii)*). Without any pragmatic constraints, 28.8% of them would incorrectly receive an ambiguous interpretation.¹³ In 77 cases (21.9%) this would even lead to a permanent ambiguity, i.e., one which is not only due to incremental processing. The pragmatic constraints previously discussed help reducing these rates to 10.3% (for ambiguous interpretations) and 8.6% (for permanent ambiguities), respectively.

Related Work

Several cognitive studies have questioned the validity of the LMF model. Reading time experiments (Gibbs, 1989; Gerrig, 1989) demonstrate that processing times to understand figurative speech are not necessarily higher than those for literal language – these evidences are not yet conclusive, however (Gibbs, 1990). Preliminary evidence has also been gathered for the claim that reasonable non-literal interpretations emerge independently from deviant literal ones (Shinjo & Myers, 1987) and that, in discourse contexts, metonymic utterances may receive higher literal ratings than without a given context (Gibbs, 1993). This parallels our experience which indicates that figurative readings are not always deviant – as many expressions are not marked by an intrasentential anomaly, though they are clearly meant figuratively.

Our model can be considered an explicit reconstruction of an informal spreading activation model described by Recanati (1995). First, the literal meaning of a constituent is activated (corresponding to our concept A.C). Then activation flows to associatively related concepts in order to determine the potential candidates for the semantic value of the constituent. This corresponds roughly to our path-finding process, which

¹²The resolution of literal nominal anaphors is not affected by our algorithm.

¹³Note that the correct interpretations are always preserved.

is more stringent, however, as it takes other concepts apart from A.C into account. All candidates, whether literal or not, are processed in parallel. The interpretation, which fits best with the discourse, is selected by preferential choice (cf. Table 1), a criterion not spelled out by Recanati.

The overwhelming majority of approaches to metonymy processing follow some variant of the LMF paradigm (Fass, 1991; Hobbs et al., 1993; Norvig, 1989; Pustejovsky, 1995; 1991; Verspoor, 1996).¹⁴ In Fass's and Norvig's work figurative interpretations are only triggered when a sortal conflict is encountered, thereby explicitly subscribing to a two-stage model. In Hobbs *et al.*'s work the literal interpretation of a sentence is regarded as a special case of metonymic resolution, thus making a unified treatment of literal and figurative language possible, in principle. Nevertheless, the special case of the literal interpretation of a sentence is strictly preferred when no sortal conflicts are encountered, regardless of discourse constraints. Pustejovsky and Verspoor both concentrate on *logical metonymy*, which apart from semantic effects has also syntactic ones. Nevertheless, type coercion is regarded as a typically *semantic* operation applied after encountering sortal conflicts (see, e.g., Pustejovsky (1991), pp.425).

In all these studies, the interaction of anaphora and metonymies is not an issue, though some of them (Norvig, 1989; Hobbs et al., 1993) consider nominal anaphora in a text understanding system framework like ours.

Consistent with our goal to avoid overgeneration, lexical semanticists like Verspoor (1996) or Godard & Jayez (1993) have proposed additional, highly sophisticated constraints on coercion they assign to the lexicon. Verspoor argues that purely pragmatic constraints cannot be sufficiently strict or explain language-specific differences. This argument cannot be defeated in principle, but we have shown that there are at least a range of pragmatic regularities, which do not require additional modelling in the lexicon but mirror general constraints at the knowledge level and are worth being considered when constraining coercion. Supplementing them with lexical and language-specific constraints is a worthwhile direction for future work.

Conclusion

We have introduced a model of metonymy resolution that treats metonymies on a par with literal interpretations. Furthermore, the mutual influence metonymies and anaphora have upon each other in noun phrases is adequately accounted for. In particular, rejecting the LMF hypothesis in favor of a parallel approach leads us to a more effective procedure.

Our empirical data have led us to several follow-up questions. In particular, the high percentage of non-anaphoric noun phrases not marked by a sortal conflict will be in the focus of our future work. Possibly these data will provide further evidence for our claim that LMF approaches are infeasible to adequately account for metonymy interpretation.

¹⁴There is one computational approach to *metaphoric* speech (Martin, 1992), however, which explicitly rejects the LMF model.

Acknowledgements. K. Markert is a member of the Graduate Program on *Human and Machine Intelligence* at Freiburg University, which is funded by DFG. We thank our colleagues, in particular Susanne Schacht and Michael Strube, who commented on earlier drafts of this paper. We also like to express our deep gratitude for the fruitful discussions we had with Bonnie Webber.

References

- Chaffin, R. (1992). The concept of a semantic relation. In A. Lehrer & E. Kittay (Eds.), *Frames, Fields and Contrasts*, pp. 253–288. Hillsdale, N.J.: Lawrence Erlbaum.
- Fass, D. C. (1991). *met**: A method for discriminating metonymy and metaphor by computer. *Computational Linguistics*, 17(1):49–90.
- Gerrig, R. (1989). The time course of sense creation. *Memory and Cognition*, 17:194–207.
- Gibbs, R. (1989). Understanding and literal meaning. *Cognitive Science*, 13:243–251.
- Gibbs, R. (1990). Comprehending figurative referential descriptions. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 16(1):56–66.
- Gibbs, R. (1993). Literal meaning and figurative language. *Discourse Processes*, 16:387–403.
- Gibbs, R. (1994). Figurative thought and figurative language. In M. A. Gernsbacher (Ed.), *Handbook of Psycholinguistics*, pp. 411–446. San Diego, Cal.: Academic Press.
- Godard, D. & J. Jayez (1993). Towards a proper treatment of coercion phenomena. In *Proc. of EACL-93*, pp. 168–177.
- Grice, P. (1989). *Studies in the Way of Words*. Cambridge, Mass.: Harvard University Press.
- Hobbs, J. R., M. E. Stickel, D. E. Appelt & P. Martin (1993). Interpretation as abduction. *Artificial Intelligence*, 63:69–142.
- Keysar, B. (1994). Discourse context effects: Metaphorical and literal interpretations. *Discourse Processes*, 18:247–269.
- Lakoff, G. (1987). *Women, Fire, and Dangerous Things. What Categories Reveal about the Mind*. Chicago, IL: Chicago University Press.
- Markert, K. & U. Hahn (1997). On the interaction of metonymies and anaphora. In *Proc. of IJCAI-97*.
- Markert, K., M. Strube & U. Hahn (1996). Inferential realization constraints on functional anaphora in the centering model. In *Proc. of CogSci-96*, pp. 609–614.
- Martin, J. (1992). Computer understanding of conventional metaphoric language. *Cognitive Science*, 16:233–270.
- Norvig, P. (1989). Marker passing as a weak method for inferencing. *Cognitive Science*, 13(4):569–620.
- Nunberg, G. (1995). Transfers of meaning. *Journal of Semantics*, 12:109–132.
- Pustejovsky, J. (1991). The generative lexicon. *Computational Linguistics*, 17(4):409–441.
- Pustejovsky, J. (1995). *The Generative Lexicon*. Cambridge, Mass.: MIT Press.
- Recanati, F. (1995). The alleged priority of literal interpretation. *Cognitive Science*, 19:207–232.
- Shinjo, M. & J. Myers (1987). The role of context in metaphor comprehension. *Journal of Memory and Language*, 26:226–241.
- Stallard, D. (1993). Two kinds of metonymy. In *Proc. of ACL-93*, pp. 87–94.
- Strube, M. & U. Hahn (1996). Functional centering. In *Proc. of ACL-96*, pp. 270–277.
- Verspoor, C. (1996). Lexical limits on the influence of context. In *Proc. of CogSci-96*, pp. 116–120.
- Woods, W. A. & J. G. Schmolze (1992). The KL-ONE family. *Computers & Mathematics with Applications*, 23(2-5):133–177.