On the Interaction of Metonymies and Anaphora

Katja Markert & Udo Hahn

Computational Linguistics Research Group
Freiburg University, Werthmannplatz 1, D-79085 Freiburg, Germany
{markert,hahn}@coling.uni-freiburg.de
http://www.coling.uni-freiburg.de

Abstract

From the analysis of naturally occurring texts we obtained evidence for the systematic interaction between nominal anaphora and metonymies. This leads us to postulate an integrated model incorporating both phenomena simultaneously. The consideration of discourse constraints for metonymy resolution allows us to challenge the commonly held view that the interpretation of metonymies should proceed from a literal-meaning-first approach. Thus, we argue for an equally balanced treatment of literal and figurative language use.

1 Introduction

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 contiguity relation r. Most computational studies concentrate on the fact that taking *A* literally often leads to *intrasentential* semantic or syntactic anomalies, which they take as a starting point for metonymy resolution. This approach disregards the textual embedding of most metonymies, which is reflected by the systematic interaction in the resolution of nominal anaphora and definite metonymic noun phrases. This is a surprising fact, since metonymy is generally regarded as a phenomenon of reference, which obviously may interact with other referential phenomena as well. Focussing on intrasentential phenomena only, almost unequivocally implies subscribing to the *literal-meaning-first (LMF) hypothesis* [Gibbs, 1989]. The most common variant of the LMF view considers the occurrence of sortal conflicts (or, deriving from that, syntactic irregularities) a necessary condition for metonymy resolution. In all other cases the literal interpretation is preferred, by default.

The subsequent discussion is intended to reveal various interaction patterns for metonymies and anaphora and the impact they have on the occurrence of sortal conflicts. We then conclude that restricting the resolution of metonymies to intrasentential analysis, as well as pursuing an LMF approach are both inadequate. Given a definite noun phrase *A,* we may distinguish between the following usage patterns:

/. *A* is *anaphoric* and ...

1. *A* is *not metonymic.*[1] 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?*

   (i) Ten minutes before the *notebook* switches off, it starts beeping,

   (ii) The clock frequency of the *computer* is reduced to 8 MHz.

   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.*[3]

   (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

[1] We restrict figurative speech in this paper to metonymies.

[2] Stallard [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.

[3] Given the original discourse context of this text fragment, *"Epson EPL-5600 "* is already known to be a laser printer.

the problem is reduced to case *b).* Example *b)* illustrates how metonymy resolution can benefit from information about possible anaphoric antecedents of *"the laser".* Firstly, the 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").* Secondly, only the information about possible antecedents can help triggering a metonymy resolution at all, as no sortal restrictions are violated.

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 *textual ellipsis* [Hahn *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 *I.3.(ii).b* shows that these criteria can also be met by anaphoric noun phrases *A* that are referential metonymies. Thus, we may get either truly ambiguous readings here or, at least, readings which cannot be distinguished without further lexical or pragmatic information (cf. lHahn and Markert, 1997]).

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.

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. They are also crucial for the distinction between predicative and referential metonymies. In addition, evidence from preceding discourse elements may override the literal-meaning-first hypothesis 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.

2 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 and Schmolze, 1992] for a survey). The concept hierarchy consists of a set of concept names *T* = {COMPUTER-SYSTEM, PRINTER,...} and a subclass relation *isaF* = {(LASER-PRINTER, PRINTER), (NOTEBOOK, COMPUTER-SYSTEM),...} $\subset \mathcal{F} \times \mathcal{F}$. The set of relation names $\mathcal{R}$ = {*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 $isa_{\mathcal{R}}$ = {*(has-laser, has-physical-part), (clock-frequency-of, property-of),...*} $\subset \mathcal{R} \times \mathcal{R}$. For every $i \in \mathcal{INST}$, the set of instances, the (direct conceptual) class of i is defined by $\bigcap_{C \in \mathcal{F}, i \in C} C$. We express this by $class(i) = C$ or by i *inst-of* C.

2.1 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., COMPUTER-SYSTEM is associated with the lexical item *"computer"*). Furthermore, *lex.r* refers to the *instance* corresponding to *lex.*[4] 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. The following basic criteria for determining well-formed paths are applied:

Given two concepts $x, y \in \mathcal{F}$, a series of conceptual relations $r_i \in \mathcal{R}$ $(i = 1,...,n)$ and concepts $c_j \in \mathcal{F}$ $(j = 0,...,n)$ ($n \in \mathbf{N}$, the natural numbers) is called *connective, iff*

- r_i is a (possibly inherited) conceptual role of c_{i-1} with $range(r_i) = c_i$ for all $i = (1,...,n)$;
- $c_0 = x \wedge (c_n\ isa_{\mathcal{F}}^*\ y \vee y\ isa_{\mathcal{F}}^*\ c_n)$, where $isa_{\mathcal{F}}^*$ denotes the reflexive and transitive closure of $isa_{\mathcal{F}}$.

In the following, a *connected conceptual path* like the one above will be denoted by $(r_1 \ldots r_n)$.

Furthermore, we require a well-formed path to be *non-cyclic,* favoring a unidirectional search in the knowledge base (cf. Table 1; s^{-1} denotes the inverse of relation s).

Non-Cyclic $((r_1 \ldots r_n))$:⇔ $\forall i, j \in \{1, \ldots, n\}, i \neq j$: $\neg \exists s \in \mathcal{R}: (r_i\ isa_{\mathcal{R}}^*\ s) \wedge (r_j\ isa_{\mathcal{R}}^*\ s^{-1})$

Table 1: Non-Cyclic Paths

Thus, we call a path $(r_1 \ldots r_n)$ *cyclic* iff it contains two relations which are inverses of each other (in this case, including inheritance of conceptual relations).

2.2 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 $\mathcal{M}$ which is labeled *"metonymic"*, and all remaining paths which are labeled *"unclassified"*. Hence, a literal path between A.C and B.C mirrors a literal interpretation of both *A* and

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

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 *l* 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 *sub*relations 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 $\mathcal{L}$. We refer to the first six of these patterns as *transitive part-whole patterns*, in short $\mathcal{T}$, and to the inverse patterns as $\mathcal{T}^{-1}$.

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. In order to determine path patterns corresponding to these metonymies consider the conceptual link between the concepts A.C and B.C, where *B* stands metonymically for an instance of the class Y. A corresponding well-formed conceptual path $p = (r_1 \ldots r_n)$ with $n \in \mathbb{N}$, $n > 1$, and $r_i \in \mathcal{R}$ $(i = 1,...,n)$ must, first, link A.C to Y via $p_1 = (r_1 \ldots r_{j-1})$ for some $j \in \{2,...,n\}$. Y is then linked to B.C via $p_2 = (r_j \ldots r_n)$.

With simple metonymies (as opposed to chains of metonymies), $p_1 \in \mathcal{L}$ must hold. The second link p_2 must express one of the metonymic relations $\mathcal{MS} = \{$*has-part, part-of, produced-by, contained-in, made-of*$\}$, depending on the specific metonymy to be handled.[5] In the case of a *producer-for-product* metonymy, e.g., $j = n$ and r_n = *produced-by* must hold. For a *part-for-whole* or *whole-for-part* metonymy, $j < n$ may be possible, as all paths in $\mathcal{T}$ and $\mathcal{T}^{-1}$ (e.g., (*has-physical-part**)) also express a single *has-part* or *part-of* relation. For notational convenience, we now consider the paths in $\mathcal{T}$ and $\mathcal{T}^{-1}$ as a single relation so that we may write (*has-physical-part**) $isa_{\mathcal{R}}$ *has-part* or (*event-feature**) $\in \mathcal{MS}$. We summarize the definition of metonymic paths in Table 2.

Considering *1.2.(ii)*, e.g., 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 path (*clock-frequency-of cpu-of motherboard-of*), a metonymic one. With the above notation, $n = 3$, $j = 2$, $p_1 = (r_1)$ = (*clock-frequency-of*) $\in \mathcal{L}$ and p_2 = (*cpu-of motherboard-of*) $isa_{\mathcal{R}}$ *part-of* holds. We deduce by this path pattern that *"computer"* denotes a *whole-for-part* metonymy for an instance of the concept CPU = *range* r_1.

Metonymic-Path $((r_1 \ldots r_n)) :\Leftrightarrow$
$(r_1 \ldots r_n) \notin \mathcal{L} \wedge \exists j \in \{2, \ldots ,n\}$:
$((r_1, r_2, \ldots, r_{j-1}) \in \mathcal{L} \wedge (r_j, r_{j+1}, \ldots, r_n) \in \mathcal{MS})$
$\vee ((r_j, r_{j+1}, \ldots, r_n) \in \mathcal{L} \wedge (r_1, r_2, \ldots, r_{j-1}) \in \mathcal{MS}^{-1})$

Table 2: Metonymic Path Patterns

Thus, example *1.2(ii)* can be handled in a straightforward way without incorporating discourse restrictions. Since only a metonymic path between CLOCK-FREQUENCY and COMPUTER-SYSTEM is found, we do not need to consider discourse constraints to resolve any ambiguities. Also, assuming a strict preference for literal over metonymic path patterns (as characteristic of LMF approaches) would not penalize metonymy resolution.[6] In contradistinction, example *1.3(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 1.3*(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 as *indefinite* noun phrases, verbs or adjectives. If a *definite* noun phrase occurs, we will incorporate information about possible anaphoric antecedents.

3 Interaction of Metonymies and Anaphora

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 and 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, $C_f(U_{i-1})$ = $[c_1.r, \ldots, c_n.r]$ 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. First, the path finder and evaluator are called with H.C and A.C as their arguments.[7] This returns a list of all well-formed paths between H.C and A.C, which are marked for their literalness or figurativeness. In the example *1.3.(ii).b*, anaphora resolution for *"the laser"* is triggered when the syntactic link between *"laser"* and *"liked"* is to be established. The path

[5] If the direction of the search is reversed (from B.C to A.C) the corresponding inverse relations must be considered. We refer to these inverse relations as $\mathcal{MS}^{-1}$ = $\{$*part-of, has-part, produces, contains, material-of*$\}$. This list of metonymic relations is, of course, incomplete and may be augmented on demand.

[6] We still cannot decide whether *"computer"* is a predicative or a referential metonymy without incorporating discourse information.

[7] In our framework, the path finder makes use of the list of possible antecedents of *A* to restrict the search and to allow additional conceptual specializations. We will neglect this influence here for the sake of brevity.

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. This brings us back to the distinctions discussed in Section 1. When A is literal and anaphoric, a literal path between A.C and H.C exists (we denote the *existence* of such a path by *literal*(H.C, A.C)) and an element is in the centering list whose class is a subconcept of A.C. Generally speaking, the existence of a literal path between D.C and A.C must be met for any item D in U_i that is related to A by a syntactic relation type σ. This relation type whether constituency-based or dependency-based, has to be selected in accordance with the underlying grammar theory. We denote the corresponding set of σ-related items D in utterance U_i by $U_i^\sigma(A)$ (cf. Table 3).

$$\begin{aligned}&\text{literal}(A) \wedge \text{anaphor}(A) \Leftrightarrow\\ &\quad (\forall D \in U_i^\sigma(A): \text{literal}(\text{D.C}, \text{A.C})) \wedge\\ &\quad \exists\, c.r \in C_f(U_{i-1}) : \mathit{class}(c.r)\ \mathit{isa}_{\mathcal{F}}{}^{*}\ \text{A.C}\end{aligned}$$

Table 3: Literal Anaphora

When A is a metonymy for an instance of the concept B via the relation $r \in \mathcal{MS}$, a metonymic path containing r between A.C and H.C exists (we denote the existence of such a path by *metonymic*(H.C, A.C)). When A is anaphoric we have to distinguish between 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 (cf. Table 4 for the corresponding generalization).

$$\begin{aligned}&\text{predicative metonymy}(A) \wedge \text{anaphor}(A) \Leftrightarrow\\ &\quad (\exists D \in U_i^\sigma(A): \text{metonymic}(\text{D.C}, \text{A.C}) \wedge \neg\, \text{literal}(\text{D.C}, \text{A.C})) \wedge\\ &\quad \exists\, c.r \in C_f(U_{i-1}): \mathit{class}(c.r)\ \mathit{isa}_{\mathcal{F}}{}^{*}\ \text{A.C}\\[1em] &\text{referential metonymy}(A) \wedge \text{anaphor}(A) \Leftrightarrow\\ &\quad (\exists D \in U_i^\sigma(A): \text{metonymic}(\text{D.C}, \text{A.C})) \wedge\\ &\quad \neg\exists\, c.r \in C_f(U_{i-1}) : \mathit{class}(c.r)\ \mathit{isa}_{\mathcal{F}}{}^{*}\ \text{A.C} \wedge\\ &\quad \exists\, c.r \in C_f(U_{i-1}): \mathit{class}(c.r)\ \mathit{isa}_{\mathcal{F}}{}^{*}\ \text{B}\end{aligned}$$

Table 4: Referential vs. Predicative Metonymy

With regard to 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 $\mathit{isa}_{\mathcal{F}}{}^{*}$ B = LASER-PRINTER holds. Thus, we end up with three interpretation hypotheses:

1. First, *"laser"* is literal (expressed by a literal path between LIKE and LASER). Then *"laser"* is not anaphoric.
2. Second, *"laser"* is a metonymy (expressed by a metonymic path between LIKE and LASER) with the following subhypotheses:
   (a) The metonymy is predicative. Then *"laser"* is not anaphoric.
   (b) The metonymy is referential. Then *"laser"* is anaphoric and the referent of *"laser", laser.r,* is given by *"Epson-EPL-5600"*.

Choosing among these readings, we propose the following ranking with respect to the ordering relation *str* in Table 5:

$$\begin{aligned}&\text{literal}(A) \wedge \text{anaphor}(A)\\ >_{str}\ &\text{predicative metonymy}(A) \wedge \text{anaphor}(A)\\ >_{str}\ &(\text{referential metonymy }(A) \wedge \text{anaphor}(A))\\ &\vee (\text{literal}(A) \wedge \neg\, \text{anaphor}(A))\\ >_{str}\ &\text{metonymy}(A) \wedge \neg\, \text{anaphor}(A)\end{aligned}$$

Table 5: Ranking Constraints

Considering our example, we encounter a case of ambiguity as fndicated in line 3 of Table 5. This results in instantiating two separate readings of sentence *1.3.(ii)b* and discarding reading *2.a)* corresponding to the last line in Table 5. We may now explain the merits and the difficulties of this ranking by looking at the consequences it has for the resolution of metonymies, anaphors and textual ellipses.

Metonymy Resolution. 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. As the literal meaning of *"computer"* allows for anaphora resolution we treat the metonymy as a predicative one (second line of Table 5). 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 1*.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.

Anaphora Resolution. Similarly, anaphora resolution proceeds as usual when a literal interpretation allows to resolve an anaphor (first line of Table 5). If several items in the center list can be considered possible antecedents, then its most highly ranked element is chosen as the most likely antecedent. Unlike LMF approaches, nominal anaphora that are referential metonymies can be resolved incrementally without the need for backtracking. So, in example *1.3.(ii).b, "the laser"* can be resolved to *"Epson EPL-5600"* in sentence *(i)* in at least one reading (third line of Table 5). Thus, different referential mechanisms are at work in the interpretation of definite noun phrases - anaphoric reference by using a more general expression *{printer* for *laser printer)* and anaphoric

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

reference by using a metonymic expression *(laser* for *laser printer)* - and they are covered by the <u>same</u> algorithm.

Textual Ellipsis Resolution. A crucial problem with our approach lies in the disjunction in line 3 in Table 5, which is a continuous source of ambiguity. Whereas it leads correctly to an ambiguity in example *I.3.(ii).h* as explained above, it also leads to the same kind of ambiguity in example *II.I.(ii)* where this is clearly not wished for. The algorithm as proposed so far does not yield any criteria to prevent this kind of overgeneration. This is not an artificial problem created by our algorithm in the first place, however. Rather, it mirrors the fact that the resolution of metonymies is not fully constrained by metonymic patterns like *part-for-whole* or *producer-for-product.* This is relatively independent of applying an LMF approach or not. In [Hahn and Markert, 1997], we propose three heuristics to further constrain coercion.

4 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 textual ellipses, all occurring in 606 utterances. Table 6 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 anaphora (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.

	direct s.c.	indirect s.c.	no s.c.	Σ
Anaphors	15 (14.6%)	27 (26.2%)	16 (15.5%)	58 (56.3%)
Oth. 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 6: Distribution of Metonymic Noun Phrases

	direct s.c.	indirect s.c.	no s.c.	Overall rate
Anaphors	93.3%	85.2%	87.5%	87.9%
Oth. def. NPs	100%	0%	0%	50%
Indef. NPs	72.7%	0%	0%	51.6%
Overall rate	84.1%	79.3%	46.7%	72.8%

Table 7: Resolution Rates for Metonymic Noun Phrases

Table 7 depicts the resolution rates for our approach.[9] The high rate of correctly resolved anaphoric metonymies (87.9%) lends some credit to our ranking in Table 5. Especially important is the fact that 85.2% of all anaphoric metonymies marked by an indirect sortal conflict and 87.5%; without a sortal conflict can be resolved by an incremental approach without backtracking - due to the timely consideration of anaphoric constraints. LMF approaches are restricted to account for only 42 anaphoric metonymies (72.4%) marked by a sortal conflict (cf. Table 6, 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%) 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 6, 58 of 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.[10] 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.

The effects our algorithm has on the resolution of textual ellipses are discussed in [Hahn and Markert, 1997].

5 Related Work

Computational approaches to metonymy resolution, up to now, have largely neglected the interdependencies between discourse processes and metonymies — regardless of whether they are based on pragmatics [Hobbs *et al.,* 1993; Norvig, 1989; Fass, 1991] or rooted in lexical semantics [Pustejovsky, 1995;Verspoor, 1996].

Nominal anaphora and metonymies have only been treated in a common framework by [Norvig, 1989] and [Hobbs *et al.,* 1993].[11] Nevertheless, neither of them considers interaction effects. Even worse, these are particularly excluded in the approach [Norvig, 1989] has chosen. He uses a spreading-activation model in which different types of inferences (e.g., metonymic and anaphoric ones) belong to different inference *classes.* While competing inferences *within one* class can be ranked and the best can be selected for resolution, ranking and selecting *across* different inference classes are prohibited. This contrasts with our approach in which the best *com-*

[9] We cannot compare our results to those from other approaches, since none of them has been empirically tested, so far. We here consider instead general problems LMF approaches are likely to face.

[10] Note that the resolution of literal nominal anaphors is not affected by our algorithm.

[11] Stallard (1993] from whom we have adapted the distinction between predicative and referential metonymies treats only intrasentential pronominal anaphora. In the lexical semantics camp, Pustejovsky [1995] mentions discourse effects but does not go into any details. Verspoor [1996] integrates lexical and pragmatic constraints concentrating on rhetorical relations between two propositions (like causality), while she does not consider anaphoric phenomena.

bination of anaphoric and metonymic inferences is chosen. Another problem with Norvig* s approach relates to the fact that insufficient results from the inference processes do not trigger additional computations in order to try yet another, perhaps not so obvious analysis. If, e.g., for reasons whatsoever no anaphoric inference is found for a definite noun phrase, Norvig's system does not attempt to find a less favored metonymic reading, which, nevertheless, would allow for a *subsequent* anaphoric interpretation.

In Hobbs et al.'s work the interpretation of a sentence (or a sequence thereof) amounts to proving its logical form. Although this allows, in principle, for the interaction between metonymies and anaphora, this problem is not yet addressed in their work. Therefore, their approach to metonymy resolution is limited to explicitly marked cases of sortal conflicts. Although literal interpretation is regarded as a special case of metonymy resolution, it is preferred when no sortal conflict is encountered, thus actually subscribing to an LMF model.

This leads us to our second point of criticism, namely that all computational approaches we know of converge on the assumption (rejected in our approach) that metonymy resolution is entirely dependent on encountering intrasentential anomalies. Usually, this amounts to the requirement that sortal conflicts have to be recognized, e.g., [Fass, 1991; Hobbs *et ai,* 1993; Norvig, 1989]; sometimes a small range of syntactic conflicts is considered as well, e.g., iGrosz *et ai,* 1987; Pustejovsky, 1995]. Thus, a metonymic interpretation of a sentence is always rejected in favor of the literal one, regardless of discourse processes, if the literal one does not lead to a deviant sentence reading. Our test set indicates, however, that about one third of the metonymies it contains are not marked by a sortal conflict at all. Hence, further relying on this violation condition precludes a significant range of metonymies from actually being resolved.

The pragmatic approaches just discussed fully subscribe to a syntax-first paradigm. The incorporation of their mechanisms for metonymy resolution into an incremental mode is precluded, as they rely on the detection of sortal conflicts. In contradistinction, our predicates in Table 3, 4 and 5 are all compatible with both a syntax-first paradigm and the incremental interaction of syntax, semantics and pragmatics.

6 Conclusion

We have introduced a model of metonymy resolution that is based on the interaction of metonymic and anaphoric resolution processes. Discourse constraints supplied by information about anaphoric antecedents enable us to treat metonymies on a par with literal interpretations. This leads us to a more effective procedure than common LMF models.

The empirical data we got from our approach - despite the fact that such an evaluation of metonymy resolution, to the best of our knowledge, has been carried out for the first time - have led us to several follow-up questions. Prime among them is a learning procedure for path patterns, one that is able to automatically determine metonymic relations starting from a given set of literal ones.

While our procedure is language-neutral, several constraints will also have to be assigned at the lexical or syntactic level of description. Combining these levels is different from the approach by [Pustejovsky, 1995], who pursues language-specific encoding in the lexicon only. As is evident, any approach to metonymy resolution is heavily dependent on the structural richness of the available knowledge sources (whether qualia structures in the lexicon or terminological world knowledge). The more elaborate the representational depth, the more sophisticated the metonymic relations that can be determined. This argument, however, is a truism and applies to any attempt at real natural language *understanding.*

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, 1992] Roger Chaffin. The concept of a semantic relation. In A. Lehrer and E.F. Kittay, editors, *Frames, Fields and Contrasts,* pages 253-288. Hillsdale, N.J.: Lawrence Erlbaum, 1992.

[Fass, 1991] Dan C. Fass. met*: A method for discriminating metonymy and metaphor by computer. *Computational Linguistics,* 17(l):49-90,1991.

[Gibbs, 1989] Raymond Gibbs. Understanding and literal meaning. *Cognitive Science,* 13:243-251, 1989.

[Grosz et al, 19871 Barbara Grosz, Douglas E. Appelt, Paul A. Martin, and Fernando C.N. Pereira. Team: An experiment in the design of transportable natural-language interfaces. *Artificial Intelligence,* 32(2): 173-243, 1987.

[Hahn and Markert, 1997] Udo Hahn and Katja Markert. In support of the equal rights movement for literal and figurative language: A parallel search and preferential choice model. In *Proc. of CogSci-97,* pages 609-614,1997.

[Hahn *et ai,* 1996] Udo Hahn, Katja Markert, and Michael Strube. A conceptual reasoning approach to textual ellipsis. In *Proc. of ECAI-96,* pages 572-576,1996.

[Hobbs *etai,* 1993] Jerry R. Hobbs, Mark E. Stickel, Douglas E. Appelt, and Paul Martin. Interpretation as abduction. *Artificial Intelligence,* 63:69-142. 1993.

[Lakoff, 1987] George Lakoff. *Women, Fire, and Dangerous Things. What Categories Reveal about the Mind.* Chicago University Press, Chicago, IL, 1987.

[Norvig, 1989] Peter Norvig. Marker passing as a weak method for inferencing. *Cognitive Science,* 13(4):569-620, 1989.

[Nunberg, 1995] Geoffrey Nunberg. Transfers of meaning. *Journal of Semantics,* 12:109-132,1995.

[Pustejovsky, 1995] James Pustejovsky. *The Generative Lexicon.* Cambridge, Mass.: MIT Press, 1995.

[Stallard, 1993] David Stallard. Two kinds of metonymy. In *Proc. ofACL-93,* pages 87-94, 1993.

[Strube and Hahn, 1996] Michael Strube and Udo Hahn. Functional centering. In *Proc. ofACL-96,* pages 270-277, 1996.

(Verspoor, 1996] Cornelia Verspoor. Lexical limits on the influence of context. In *Proc. of CogSci-96,* pages 116-120,1996.

[Woods and Schmolze, 1992] William A. Woods and James G. Schmolze. The KL-ONE family. *Computers & Mathematics with Applications,* 23(2-5): 133-177,1992.