

TWO KINDS OF METONYMY

David Stallard
BBN Systems and Technologies
70 Fawcett Street
Cambridge, MA 02138, USA
Internet: stallard@bbn.com

ABSTRACT

We propose a distinction between two kinds of metonymy: “referential” metonymy, in which the referent of an NP is shifted, and “predicative” metonymy, in which the referent of the NP is unchanged and the argument place of the predicate is shifted instead. Examples are, respectively, “The hamburger is waiting for his check” and “Which airlines fly from Boston to Denver”. We also show that complications arise for both types of metonymy when multiple coercing predicates are considered. Finally, we present implemented algorithms handling these complexities that generate both types of metonymic reading, as well as criteria for choosing one type of metonymic reading over another.

1 INTRODUCTION

The phenomenon of semantic coercion, or “metonymy”, is quite a common one in natural language. In metonymy, the actual argument of a predicate is not the literal argument, but is instead implicit and related to the literal argument through an implicit binary relation. For example, in the following utterances, taken from Lakoff and Johnson (1980):

(1) The ham sandwich is waiting for his check

(2) Nixon bombed Hanoi

it is not literally the ham sandwich which is doing the waiting, but rather the person who ordered it, and not literally Nixon who is doing the bombing, but rather the pilots under his command. The noun phrase - “The ham sandwich”, “Nixon” - is said to be “coerced” through an implicit binary relation to a related object which is the actual argument of the predicate.

Perhaps the most familiar definition of metonymy from the literature is that it is a figure of speech in which the speaker is “using one entity to refer to another that is related to it” (Lakoff and Johnson,1980). This definition is quite commonly held in one form

or another. (For example, see (Fass,1991), where it is directly quoted; also similar definitions in (Pustejovsky,1991), (Hobbs,1988)). But what does it really mean? Does it mean that the coerced noun phrase is actually an indirect reference to an object different from its literal referent?

If so, then we might expect other linguistic data to support this. For example, we might expect subsequent anaphora to agree with the “real” referent. And indeed, in the following dialogue the intra-sentential pronoun “his” and the extra-sentential “he” both agree with the indirect reference to the customer, not the the literal sandwich:

(3) The ham sandwich is waiting for his check
He is getting a little impatient

But compare the dialogues

(4) Nixon bombed Hanoi.
He wanted to force the Communists to negotiate

(4') Nixon bombed Hanoi
*They sang all the way back to Saigon

The dialogue (4) is quite natural, while in (4'), the use of “they” to refer to the bomber crews seems ruled out - the reverse of what the indirect reference view would predict.

A second problem with the indirect reference view is found in certain performative contexts, such as wh-questions and imperatives, in which the referent of a particular NP is sought by the speaker. If this NP is metonymically coerced, we could expect the correct response to the utterance to be the indirect reference. Consider, the following examples, which are actual utterances collected for the DARPA ATIS domain (MADCOW,1992), a database question-answering domain about commercial air flights:

(5) Which wide-body jets serve dinner?

(6) Which airlines fly from Boston to Denver?

In ATIS, only flights “fly” or “serve meals” and thus both sentences can only be understood metonymically. In (5), it is not the jets which serve dinner but the flights on the jets, and one plausible construal is indeed that “wide-body jets” is really a reference to flights on wide-body jets, and the interpretation of the sentence is a request to display the set of these flights. This would agree with the indirect reference view.

In (6), however, the only possible construal seems to be that a set of airlines - the airlines offering flights from Boston to Denver - is being sought. To respond to this request with the set of flights from Boston to Denver would clearly be absurd.

We propose a distinction, motivated by such examples, between two kinds of metonymy, which we term *referential* and *predicative*. In referential metonymy, the metonymic noun phrase does indeed have an intended referent related to but different from its literal meaning. An example is the noun phrase “the ham sandwich” in (1) above, where the actual and intended referent is to a related object - the person who ordered the sandwich. In predicative metonymy, however, the actual and intended referent of the noun phrase is just the literal one, and it is more accurate to say that the predicate is coerced (though as we show later, this is itself a simplification). An example of predicative metonymy is (6) above.

We also show how both types of metonymy are complicated by the presence of multiple predicates that require the same coercion of an NP. We present algorithms for generating the two types of metonymic reading that cope with these complexities. Finally, we present criteria for determining a preference for one type of metonymic reading over another. (We do not, however, deal in this paper with the question of how to determine which relations to use for coercion, viewing this as a separate problem.)

The examples throughout are taken from the ATIS domain, a domain with a pre-established formal conceptual system of categories and relations that utterances must be mapped onto. The algorithms presented are implemented in the DELPHI system (Bobrow et al, 1991), which has been ported to that domain and formally evaluated in it.

The remainder of the paper is organized into the following sections:

Section 2, the next section, formalizes the distinction between referential and predicative metonymy by giving logical form readings for each, and shows how both types of metonymy are globally complicated when

multiple coercing predicates are considered.

Section 3 gives an algorithm for generating both types of metonymic readings in semantic interpretation that handles these global complications.

Section 4 gives criteria for picking one type of reading over another

Finally, section 5 compares our work to previous work on metonymy.

2 METONYMY AND LOGICAL FORM

In this section we sharpen and formalize our notion of referential and predicative metonymy by giving logical form readings for the different cases.

The logical language we use has sortal quantifiers, with a special quantifier “WH”. A wh question is the treated as:

(7) (wh x S (and (P1 x) (P2 x)))

which is interpreted as a request to display all members of S (the semantic class of the wh-np) which satisfy both P1 (the modifiers of the wh-np) and P2 (the predicate of the clause). A labeled-argument notation is used for clause semantics.

Now, let us return to the examples of the previous section. In (5), the referential metonymic reading of the sentence in which *flights* are sought that serve dinner and are on wide-body jets is expressed as:

(8) Which wide-body jets serve dinner?

(wh x flights
 (and (exists y jets
 (and (aircraft-of x y)
 (wide-body y)))
 (serve flight-of x
 meal-of dinner)))

where coercion relation is AIRCRAFT-OF, mapping between flights and the aircraft they are on.

Compare this with the reading for (6), in which airlines and not flights are sought:

(9) Which airlines fly from Boston to Denver?

(wh x airlines
 (exists y flights
 (and (airline-of y x)
 (fly flight-of y
 orig-of Boston
 dest-of Denver))))

The readings of the referential (8) and the predicative (9) are in a sense inside-out versions of each other. Both have an interpolated quantifier for FLIGHTS that is not explicitly present in the utterance but in (8) the interpolated is on the outside and is the WH-thing displayed whereas in (9) the interpolated quantifier is on the inside, and is merely part of the description of what is to be displayed. This, in logical terms, is the crux of the referential/predicative distinction.

Predicative metonymy can be loosely thought of as coercion of a predicate argument place, rather than of the argument NP itself. It may therefore seem attractive to try to formalize this in a directly compositional way through some device such as lambda-abstraction. If P is the predicate, R the binary relation of coercion, and *i* the argument-number of P to be coerced, the coerced version of P might then be defined in such a view as:

$$(\lambda(x_1, \dots, x_i, \dots, x_n) (\text{exists } y (\text{domain } R) (\text{and } (R \text{ } y \text{ } x_i) (P \text{ } x_1, \dots, y, \dots, x_n))))$$

This is a predicate just like P, except extended by the relation R in its *i*'th argument place to take an object in the range of R. Metonymic extension of the predicate would then be an essentially compositional, local process, taking place at the juncture of predicate and argument and not affecting interpretation elsewhere.

Unfortunately, such a treatment turns out to give the wrong interpretation when multiple predicates requiring the same coercion are present. Consider:

- (10a) Which airlines flying from Boston to Denver leave at 3 pm?
 (10b) Show airlines flying from Boston to Denver leaving at 3 pm

Both examples are predicative metonymic utterances. Airlines neither “fly” nor “leave”; flights do these, so both the main verb and the relative clause modifier predicates require airline-to-flight coercions. If the lambda-abstraction scheme is right each predicate-application could be dealt with separately.

Yet the following reading for 10a, which would result from the application of the lambda-abstraction scheme to the two predicates, is emphatically *not* the correct reading:

- (11) (wh x airlines
 (and (exists y flights
 (airline-of y x)
 (fly flight-of y
 orig-of Boston
 dest-of Denver))
 (exists y' flights
 (airline-of y' x)
 (leave flight-of y'
 time-of (3 pm))))))

These truth-conditions are too weak, as they allow airlines that have a Boston to Denver flight at any time, so long as they have another (possibly different) flight at 3 pm to any place. The proper reading is instead:

- (12) (wh x airlines
 (exists y flights
 (and (airline-of y x)
 (fly flight-of y
 orig-of Boston
 dest-of Denver)
 (leave flight-of y
 time-of (3 pm))))))

in which the airline is related to a single flight description that has all the desired properties.

Note that the issue here is not that one predicate is internal to the NP and the other external to it. The same problem arises with whatever combination of internal and external predicates. In 10b, for example, both predicates are internal to the NP but if the two coercions are carried out separately the same erroneous truth-conditions will result, in which the AIRLINE is related to two different FLIGHT descriptions instead of one.

Nor is the “single-interpolation” requirement related specifically to the referential/predicative distinction. If we modify one of our referential examples to include multiple coercing predicates, as below:

The ham sandwich at table 12 is impatient

we see that a correct reading would still require that the ham sandwich be related to one and only one interpolated description of a person that ordered the sandwich, is seated at table 12, and is impatient.

That fact that multiple coercions of the same NP, whether internal or external to it, cannot be carried out separately means that the phenomenon of metonymy takes on a decisively global character, one which is as much akin to quantifier scoping as it is to compositional

semantic interpretation. As we shall see in the next section, the quantifier scoping stage of processing is exactly where we locate the solution to these problems.

3 GENERATING THE READINGS

We now show how referential and predicative metonymic readings are generated, and how the requirement of a single interpolate for multiple coercions is enforced.

3.1 Input Representation

A two-stage mechanism of semantic interpretation is used in the DELPHI system. In the first stage, an initial predicate-argument level of semantic representation is produced, with quantifiers in place. In the second, a fully quantified logical form is generated, in which quantifiers are pulled out of the predicate-argument representation and placed in their proper relative scope. It is in this second stage that the referential/predicative distinction is made.

The first stage of semantic interpretation has been described elsewhere in (Bobrow et al,1991), and we do not discuss it here except to describe its output, which forms the input to the quantification stage. This output is a tree of whose nodes are phrasal representation objects. Each of these phrasal representation objects has a *head* and a set of *bindings*. The head includes semantic type information (as well as other information such as subcategorization etc.), while the bindings represent the semantic effects of modifiers on the head. Each binding has four parts:

1. the modifier grammatical relation
2. the modifier semantic relation
3. the filler of this semantic relation
4. a binary coercion relation

The following is the top-level phrasal representation for “Which airlines fly from Boston to Denver?”:

CLAUSE:
head: fly
subject: flight-of, (wh airlines), airline-of
pp: orig-of, Boston, identity
pp: dest-of, Denver, identity

This representation has three bindings: a SUBJECT and two PP-complements. In the two PP bindings, the ranges of the modifier semantic relations ORIG-OF and

and DEST-OF are both CITY, which agrees with the explicit fillers BOSTON and DENVER. Thus, in these bindings no coercion is needed and the coercion relation is just IDENTITY. But in the SUBJECT binding, the range of the modifier relation FLIGHT-OF is FLIGHT and the explicit filler is an AIRLINE. Here, the coercion relation AIRLINE-OF is required to bridge the gap between FLIGHT and AIRLINE.

NP semantic representations have the same structure, plus a quantifier. Here is the representation for “which airline”

NP:
head: AIRLINE
quant: wh

We refer to the constituent modifier bindings of the NP itself as its “internal” bindings. In this particular example, there are no internal modifiers and thus no internal bindings. When an NP is a constituent of a clause (or is the object of a PP which is), we call the binding in which the NP occurs its “external” binding.

Semantic representations of this kind are neutral not only with respect to quantifier scoping, but to the distinction between predicative and referential metonymy as well. From the standpoint of the predicate, one can think of the coercion relation as extending the given argument place of the predicate to take an argument of a different type. From the standpoint of the NP argument, on the other hand, the coercion can be viewed as mapping the NP in the “reverse” direction of the relation, from range AIRLINE to domain FLIGHT instead of from domain to range.

3.2 Algorithm

The alternative metonymic readings are generated from these semantic representations as part of the quantifier scoping pass. There are two steps.

Step 1, carried out before quantification begins, is to walk the phrasal representation tree and build a “coercion table” relating each nominal head N to the set of coercion relations on it:

R_c – the coercion relation of N’s external binding
 R_i – the coercion relations of N’s internal bindings

As a technical convenience, IDENTITY relations in the R_c, R_i are subscripted with the semantic type restriction T of the binding in which they occur. This type restriction is simply the range of the semantic modifier relation in the binding.

Step 2 is to pull the quantifiers out of this structure and into their proper places in a complete formula. For an NP with a non-IDENTITY entry in the coercion relation, alternative WFF-generating schemas are used to generate the alternative referential and predicative metonymic readings.

In what follows, let N be the noun phrase under consideration and let Q be its quantifier and S its sort. Let R be any relation which is not IDENTITY and which is one of the coercion relations associated with N in the table - whether R_e or one of the R_i .

Now, consider all the internal bindings of N which have R as their coercion relation. Let the M_R and F_R be respectively the semantic modifier relations and corresponding arguments of these bindings. Let the M_O , F_O and R_O be, respectively, the semantic modifier relations, arguments and coercion relations of bindings which do *not* have R as their coercion relation.

The operation of reading generation is to pick a non-IDENTITY R from N's table, and apply the two schemas. To generate the predicative reading, the following schema is used:

(Q x S
 {(M_O x F_O)}
 (exists y (domain R)
 (and (R y x)
 {(M_R y F_R)})
 ***)

The token “***” indicates the open slot for the matrix predicate of the clause, and the brackets “{”, “}” are shorthand for conjoined iteration over the subscripted items within.

Note that the interpolated EXISTS quantifier has scope over the matrix formula of the clause, so it will govern any external R coercion for N. Furthermore, because all the R coercions are gathered together in this scheme, the same quantifier will govern any R coercions which are internal to the NP. This fulfills the requirement of the previous section: that there be one and only one quantifier for a given coercion, even when that coercion is needed both by internal modifier relations and by the external clause in which the noun phrase is contained.

Use of the schema for our example above generates the interpretation:

Which airlines fly from Boston to Denver?

(wh x AIRLINE
 (exists y FLIGHT

(AIRLINE-OF y x)
 (FLY flight-of y
 orig-of Boston
 dest-of Denver)))

as desired.

In order to enforce the restriction that subsequent anaphora resolve to the literal AIRLINE and not the interpolated FLIGHT (and, similarly, to “Nixon” instead of the pilots in our earlier example) we add a diacritic to the interpolated quantifier ‘(exists y FLIGHT ...)’ that forbids the discourse component from resolving an anaphor to this quantified description.

The referential metonymic reading is generated by a different schema. In order to use this schema, the following condition must hold:

$(R_e = R) \vee (R_e = IDENTITY_T \wedge (domain R) \subseteq T)$

This condition ensures that a semantically ill-formed expression will not result and simply requires that the type requirement of the external binding of the NP to be referentially coerced agrees with the coerced version. Either the coercion must be dictated by the external binding itself, or the external binding's type requirement must be loose enough to accept the coerced version (as in the case of a loosely-typed predicate like “show”).

If these conditions hold, then the following schema can be used to produce the referential reading:

(Q x (domain R)
 (and (exists y S (and {(M_R y F_R)} (R x y))
 {(M_O x F_O)}))
 ***)

Use of the schema generates the following reading for our example:

Which wide-body jets serve dinner?

(wh x FLIGHT
 (exists y JET
 (and (WIDE-BODY y)
 (AIRCRAFT-OF x y)
 (SERVE flight-of x
 meal-of DINNER))))

In principle, of course, a given NP's entry in the coercion table can have more than one distinct non-IDENTITY coercion relation. Obviously in such a case there can be at most one referential coercion of the NP.

All other coercions to different semantics types must then be predicative. In the case of multiple predicative coercions, the predicative schema is simply iterated. We arbitrarily disallow chains of coercions (“double-shifting”), though these in principle could be accommodated.

4 DETERMINING THE CORRECT READING

Thus far we have argued for different types of metonymic reading and shown how to generate them, but have not given any indication of when a given type of reading is to be preferred. How do we know, for example, that the predicative reading and not the referential is correct in (6) “Which airlines fly from Boston to Denver”?

A few criteria are fairly obvious. One we have already seen in the previous section: the external-binding agreement condition on applying the referential metonymy schema. If an NP’s external semantic context agrees with its literal referent, but not its referentially coerced version, then referential metonymy is ruled out for that NP.

A somewhat broader notion of external semantic context is found in intra-sentential anaphora:

The ham sandwich is waiting for HIS check

Which airline flies to ITS headquarters city?

Clearly, we would prefer any intra-sentential anaphora to agree with the “real” referent of the NP. In the first sentence above, the pronoun “his” cannot agree with the literal referent, but can agree with the metonymically interpolated PERSON, and so provides evidence for the referential reading. In the second sentence, the pronoun “its” cannot agree in number with the interpolated set of FLIGHTS, but can agree with the singular “airline”, and so provides evidence for the predicative reading.

Neither of these two criteria addresses example (6), however. Our hypothesis is that the real distinction being made here is pragmatic. An important principle of language use (essentially part of the Gricean Maxim of Quantity (Grice,1975)) is that a cooperative speaker will avoid adding a part of a description which self-evidently adds no constraint to the set of things being described. This is the reason why such pleonasms as “female woman” sound odd to us, and are not normally

uttered. In this light, the referential reading of the sentence above:

(wh x FLIGHT
 (exists y AIRLINE
 (AIRLINE-OF x y))
 (FLY flight-of x
 orig-of Boston
 dest-of Denver))

has a completely redundant component, since every flight is on some airline. Yet this redundant component is precisely the one introduced to handle the coercion! Encoding the reference in this way has no utility: one might as well have said “which flights” to begin with.

We can formalize this principle as follows. Let R be the coercion relation and let S be the literal NP referent-set. Then the referential coercion of the NP can be written as the pairing (R,S), which describes a property on the domain of R that picks out just the subset of the domain of R that is obtained by mapping S back into the domain in the “reverse” direction of R. Such a property is considered vacuous if it provides no constraint on the domain, or in other words if:

R is a total relation and S = (RANGE R)

holds. A total, or “into”, relation is one which maps every element of its domain to at least one element of its range. Since every flight in ATIS is on an airline, AIRLINE-OF is a total relation, and AIRLINE is its range, so a referential metonymy is clearly vacuous in this case.

In contrast, the relation AIRCRAFT-OF is total, but “wide-body jet” is a proper subclass of its range (AIRCRAFT), so this condition does not hold for “What wide body jets serve dinner?” and referential metonymy is allowed for it.

Similar pragmatic considerations can be applied to rule out predicative metonymy in some cases. If a metonymically extended predicate provides no constraint on the NP, then predicative metonymy is the less likely reading. Consider again our referential example, “What wide-body jets serve dinner”. If this is taken predicatively, it would have as its logical form:

(wh x jet
 (and (wide-body x)
 (exists y flight
 (and (aircraft y x)
 (serve flight-of y
 meal-of dinner))))))

The class AIRCRAFT in ATIS is really the set of aircraft-types, and the same aircraft-type is typically used by a large number of flights with nothing particularly in common. It therefore seems unlikely that the property “(used on flights)that serve dinner” offers any constraint on the class AIRCRAFT: in other words, that being a particular type of aircraft and being used by a flight that serves dinner are correlated in any way. This particular judgment, however, is based on human knowledge and plausibility, and is difficult to formalize given the current state of the art in knowledge representation.

We have proposed a number of possible theoretical criteria for choosing between predicative and referential metonymy. It is of some interest, therefore, to compare the relative occurrences of predicative and referential metonymy in actual data. Our study of a large (> 5000 sentence) corpus of naturally collected ATIS data shows that predicative metonymy is very common. Noun phrases headed by “fare”, “airline” and “ticket” frequently appear in positions that require a flight argument. Yet it is clear, both from the meaning of the utterance, and from the judgements of independent annotators who pair these sentences with “correct” responses for NL system evaluation, that fares and airlines are being talked about in such cases, and not flights.

Indeed, our experiments have shown that allowing predicative metonymic coercion when evaluating DELPHI against this corpus leads to a 27% decrease in weighted error over not allowing it. This is very substantial difference indeed, and testifies to the importance of the metonymy phenomenon in actual data.

As for the referential type of metonymy, we have found only a few cases of it in this corpus. We hypothesize that the reason for this is that referential metonymy, involving as it does an encoding of a reference in terms of a categorially different thing, is a more marked and unusual event in psychological terms. Predicative metonymy, on the other hand, involves no such operation, merely the convenient making-way of a predicate for a non-standard but related argument. For this reason, our work prefers predicative metonymy as the default choice in processing when no other evidence is present.

5 COMPARISON WITH PREVIOUS WORK, CONCLUSIONS

We have argued for a distinction between two types of metonymic reading, and have given evidence that

metonymy must ultimately be treated as a global phenomenon over the sentence, part of which belongs with quantificational considerations and part with local compositional interpretation. We have shown how pragmatic considerations of language use can influence which reading is preferred.

The referential/predicative distinction is not observed in most of the writing on metonymy, which is either not formal and computational in nature (Lakoff and Johnson,1980), or is oriented towards different types of systems and computational concerns. Hobbs (1987,1988), for instance, discusses metonymy along with a number of other “local pragmatic” issues (nominal compounds, etc.), but this work is done in the context of a message-processing and not a question-answering system, so many of the issues we have discussed (wh-questions, etc.) simply do not arise there.

Something like the referential/predicative distinction does seem to be present, however, in the work of a few other authors. For example, Fass (1991) speaks of what he calls the “source” or the “target” of a metonymy being alternatively substituted for. His sentence representations are not done in a formal logical framework, however, so it is difficult to tell if the ambiguity has a referential or truth-conditional consequence.

Closer to our work is that of Pustejovsky (1991). He defines a notion he calls “logical metonymy” which seems quite close to our notion of predicative metonymy. In a sentence like “Mary enjoyed the book”, logical metonymy changes the type of the verb “enjoy” to take an object like “book” which is not an event but which is related to one (the reading of the book). As we have shown in Section 3, however, the single-interpolation requirement for multiple coercing predicates poses a technical problem for a verb type-changing view which only looks at the given verb and argument by themselves. Our work has demonstrated that a correct account of metonymic coercion must, in the most general case, involve considerations that are global over the whole utterance interpretation.

Our work has also demonstrated an important interaction between appropriateness of metonymic readings and the Gricean Maxim of Quantity. To our knowledge, no other work has done this. Finally, our work differs from previous work in the area by having been carried out in an environment of objective evaluation, an environment whose rigors have pushed us towards many of the insights presented here.

6 Acknowledgments

Volume 17, Number 1
March 1991

The work reported here was supported by the Advanced Research Projects Agency and was monitored by the Office of Naval Research under Contract No. N00014-92-C-0035. The views and conclusions contained in this document are those of the author and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the United States Government.

I would like to thank James Pustejovsky and Rusty Bobrow for valuable comments and discussion.

References

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

MADCOW Committee (Hirschman, Lynette et al) *Multi-Site Data Collection for a Spoken Language Corpus*
in Proceedings Speech and Natural Language Workshop
February 1992

Grice, H. P. *Logic and Conversation* in P. Cole and J. L. Morgan, *Speech Acts*,
New York: Academic Press, 1975

Pustejovsky, James *The Generative Lexicon*
Computational Linguistics
Volume 17, Number 4
December 1991

Hobbs, Jerry R. and Stickel, Mark *Interpretation as Abduction*
in Proceedings of the 26th ACL
June 1988

Bobrow, R., Ingria, R. and Stallard, D.
The Mapping Unit Approach to Subcategorization
in Proceedings Speech and Natural Language Workshop
February 1991

Hobbs, Jerry R., and Martin, Paul *Local Pragmatics*
in Proceedings, 10th International Joint Conference on Artificial
Intelligence (IJCAI-87). Milan, Italy

Fass, Dan *met*: A Method for Discriminating Metonymy and Metaphor by Computer*
Computational Linguistics