

Lexical Limits on the Influence of Context

Cornelia M. Verspoor

Centre for Cognitive Science

University of Edinburgh

2 Buccleuch Place

Edinburgh EH8 9LW

Scotland, UK

kversp@cogsci.ed.ac.uk

Abstract

This paper introduces an approach to modeling the interpretation of semantically underspecified logical metonymies, such as *John began the book*. A distinctive feature of the theory presented is its emphasis on accounting for their behavior in discourse contexts. The approach depends on the definition of a pragmatic component which interacts in the appropriate manner with lexicosyntactic information to establish the coherence of a discourse. The infelicity of certain logical metonymy constructions in some discourses is shown to stem from the non-default nature of the lexicosyntactically determined interpretation for such constructions. The extent of the influence of contextual information from the discourse on the interpretation of logical metonymies is therefore constrained by the lexical properties of the constituents of the metonymies. Contextually-cued interpretations are shown to be unattainable when infeasible lexical information conflicts with these interpretations.

Introduction

It is sometimes assumed that with a sufficiently strong context any interpretation of a semantically underspecified sentence can be coerced. If we consider, for example, the sentence (1a) independent of a context we are likely to interpret it as (1b) or possibly (1c). If we then insert the prior context *John is my pet goat; he loves eating things* we suddenly prefer the interpretation in (1d).

- (1) a. John enjoyed the book.
b. John enjoyed reading the book.
c. John enjoyed writing the book.
d. John enjoyed eating the book.

This strong influence from context does not, however, operate without limits. There are underspecified constructions which, perhaps surprisingly, appear to disallow any interpretation other than the conventional, context-independent possibilities. These constructions are infelicitous in a context which would seem to require an interpretation at variance with these conventions.

In this paper, I will focus on one such construction, that of *logical metonymies*, for which more meaning than is directly attributable to sentential components arises. Logical metonymies generally occur when a verb has alternate syntactic complement forms, with only a single semantic interpretation for all forms. For example, the sentences in (1a-b) and (2) can be viewed as expressing the same meaning, although in the (a) sentences no reading event is explicitly mentioned.

- 2) a. John began the book.

- b. John began reading/to read the book.

The systematic syntactic ambiguity of aspectual verbs and verbs like *enjoy* has been handled in existing approaches to logical metonymy via an operation of *type coercion* (either triggered by constraints on the logical type of the verbal complement, e.g. Pustejovsky (1991, 1995), or internalized in the verb semantics, e.g. Copestake and Briscoe (1995)) such that the logical forms for each verb+complement form will be identical. The coercion which must occur to get the appropriate readings of (1-2a) requires that a missing element of meaning, i.e. the reading event, be introduced. Pustejovsky (1991) has proposed that this element corresponds to one of the roles in the lexical semantic structure of the noun in the complement, the *qualia structure*. Type coercion looks to the qualia structure for an element of meaning with the logical type required by the semantics of the verb.

It could also be argued that the missing element is filled in from context via pragmatic processing rather than from a richly structured lexical representation. There are strong arguments, however, for lexical specification of some aspects of metonymy (see e.g. Copestake (1992) for discussion). An additional argument, which will be provided in this paper, is that a purely pragmatic approach would fail to constrain the possible interpretations of logical metonymies and therefore would fail to account for the full range of data. Consider the discourse in (3), for example. The context clearly cues an interpretation of the sentence (3b) of *John will begin destroying the books tomorrow*. However, this sentence is infelicitous in this discourse. In contrast, (3c) is felicitous in the discourse and has the expected interpretation.

- (3) a. John will be audited by the IRS, so he has been destroying things which might incriminate him. He has destroyed the files and the computer disks.

b.* He will begin the books tomorrow.

c. He will begin on/with the books tomorrow.

A purely pragmatic approach would fail to constrain the possible interpretations of logical metonymies, while a purely lexical approach fails to accommodate the potential contextual influence on these interpretations (as required to explain the possibility of (1a) being interpreted as (1d)). A combination of the two approaches is necessary to explain the range of logical metonymy data. In this paper, I will show that information derived in the lexicon can be used to constrain the possible interpretations of a phrase in such a way that even a strong context cannot override the lexical specifications. Furthermore, I will argue that these specifications could not be

relegated to the pragmatic component without a reduction in the generality of the treatment of logical metonymy phenomena. Thus, lexical specification of conventions are necessary, and the pragmatic component must be able to utilize the information coming from the lexicon in the appropriate ways.

Interpretation of Logical Metonymies

Lexical Approaches

As introduced above, the lexically-based explanations of logical metonymy depend on the representation of core lexical semantic information for nominals in the form of qualia structure. This structure specifies four essential aspects of a word's meaning, described with respect to the denotation of the word (Pustejovsky, 1991, 1995) — CONSTITUTIVE: the relation between the object and its constituent parts; FORMAL: that which distinguishes object within a larger domain; TELIC: the function of the object, what is done with it; and AGENTIVE: how the object came into being.

The roles in qualia structure relevant to logical metonymy are the telic and agentive roles, as these two roles will specify eventualities involving the denotation of the noun in the NP complement. For example, the predicate in the telic role of *book* is *reading*, while the predicate in its agentive role is *writing*. Most existing approaches assume that all concrete nouns always have both the telic and agentive roles specified, and thus the eventualities specified there are always available to the process of type coercion when establishing the interpretation of a logical metonymy structure. No other interpretations will be available, correctly ruling out the specified interpretations of the sentences (4), since the desired eventualities do not fill a role in the qualia structure of the nominal objects.

(4) John began the stone (**moving*) / the book (**destroying*) / the desert (**crossing*)

Type coercion, however, must be constrained in some way, as there are interpretations of metonymies predicted on the basis of the eventualities in qualia structure which are actually ungrammatical, such as those in (5). The constraints which must be added to a qualia structure approach to type coercion to rule out such examples have been discussed by Godard and Jayez (1993) and Pustejovsky and Bouillon (1995).

(5) John began the highway (**driving on*) / the dictionary (**consulting*)

The constraint which Godard and Jayez (1993) propose to rule out these cases is that the reconstructed event should be a kind of *modification* (expressing an intuition that the object usually comes into being, is consumed, or undergoes a change of state). Pustejovsky and Bouillon (1995) develop constraints on type coercion in terms of the aspectual properties of the reconstructed event. Their account relies on a structured representation of events, in which subevents are represented and the "focus" of the event is marked as the *head* of the event structure. Left-headed structures correspond to accomplishments, while right-headed structures correspond to achievements. The aspectual constraint on type coercion, then, is that the complement of *begin* must be a left-headed TRANSITION. Sentences like (5) are therefore ruled out because *driving on the highway*, etc. are activities without a definite endpoint, rather than left-headed transitions.

These constraints, however, do not rule out all implausible metonymies, and rule out some plausible ones. The sentences in (6) should be ruled out on the specified interpretations, despite the associated events all being left-headed transitions specified in qualia structure. In contrast, those in (7) would incorrectly be ruled out since the interpretations convey activities.

(6) John began the film (**watching*) / the nails (**hammering in*) / the door (**opening, *walking through*)

(7) John began daycare at his mom's work (*attending*) / the violin when he was five (*playing*) / acupuncture in April and homeopathy in August (*undergoing*)

The solution to this argued for in Verspoor (1996) rejects the assumption that all nouns have telic roles specified in qualia structure and rejects the effectiveness of the proposed aspectual constraints in capturing the range of logical metonymy data, while accepting that the context-independent natural interpretation of logical metonymies does appear to correspond to either the telic or agentive role of the noun in the NP complement. It is therefore proposed in that paper that not all artifacts have a conventionalized telic event in the qualia structure. Those that do not have this conventionalized event are infelicitous in logical metonymic constructions. Examples such as (5) and (6) are ruled out because the relevant eventualities are not specified in the qualia structure, while they are for (7).

Combined Lexical and Pragmatic Approaches

Lascarides and Copestake (1995) (L&C) extend the lexical approaches to logical metonymy to develop a system which takes into account the potential influence of context on the reconstruction of an event in a coercive environment. They utilize the idea that lexical defaults — defaults specified in qualia structure — persist beyond the lexicon into the pragmatic component, and are therefore potentially overridden by default pragmatic information.

L&C formalize their approach in a unification-based framework, with a theory of lexical structure in which the lexical entries are typed default feature structures (TDFSs), and which utilizes persistent default unification (PDU) (Lascarides *et al* 1996). The lexicon is hierarchically ordered (as described in e.g. Copestake (1992)). They adopt Pustejovsky's notion of qualia structure, which provides lexical defaults. For example, the telic role of *book* is *read* by default. Other aspects of lexical representation follow the lexical representation language (LRL) (Copestake, 1993b).

In the pragmatic component DICE (Lascarides and Asher, 1991), L&C propose two axioms, i) *Defaults Survive*: lexical generalisations normally apply in a discourse context and ii) *Discourse Wins*: conflicting discourse information wins over lexical defaults. These axioms together can be used to explain why (8a) has the interpretation (8b) rather than (8c), even though the telic role of *book* is *read*.

(8) a. My goat eats anything. He really enjoyed your book.
b. The goat enjoyed eating your book.
c. The goat enjoyed reading your book.

Since L&C build on the Pustejovsky approach to logical metonymy, they adopt the assumption of full representation of

qualia structure and rely on constraints applied prior to pragmatic processing to rule out implausible metonymies. The approach therefore suffers from the same over- and under-generation as purely lexical approaches, but it is fully compatible with the Verspoor (1996) view on what information is lexically specified. I now turn to the interaction of that lexical information with the pragmatic component, and show how this can be used to account for incoherent discourses such as that in (3a,b).

Lexical Constraints Interacting with Context

An Example

The discourse in (9) exemplifies a contrast between the behavior of *begin* (or any aspectual verb which may be substituted) and other verbs. The sentence (9c(i)) is infelicitous as a continuation of the discourse (9a,b), while the sentences (9c(ii-iii)) are not. This parallels the distinction in (3). Assuming that default interpretations for *begin (on) your book* and *enjoy your book* are predicted from the lexicon, the example suggests that the default interpretation of *begin+NP* cannot be overridden by contextual cues for its interpretation, while the default interpretations of *begin on+NP* and *enjoy+NP* can be.

- (9) a. My goat went nuts last night.
 b. He ate everything in his cage.
 c. i.*He began your book at 9pm.
 ii. He began on your book at 9pm.
 iii. He particularly enjoyed your book.

This contrast was modeled in Verspoor (1996) by having a syntactic rule associated with the structure *aspectual verb+NP* to trigger the operation DefFill (defined in Lascarides *et al* 1996). This operation converts the typed default feature structure (TDFS) representing the lexically specified default interpretation of the *aspectual verb+NP* phrase to a non-default typed feature structure (TFS) prior to pragmatic processing, thus restricting the interpretations which persist into the pragmatic component to those that are lexically specified. This operation is not triggered in the cases of *begin on* or *enjoy* or similar verbs, and thus in these cases the lexically specified default interpretation remains default and can be overridden by contextual information via the axiom *Discourse Wins*.

It was claimed in Verspoor (1996) that the conflict in pragmatics between the non-default information coming from the lexical/grammatical components and the interpretation suggested by context can result in a judgement of the sentence in question as infelicitous, but the process underlying this judgement was not explained. I will outline it here, showing how the judgement crucially depends on the non-default nature of the lexically specified interpretation for *begin the book*, which leads to an inability to connect the sentence coherently to the preceding discourse.

The Analysis

DICE is a theory which allows us to compute rhetorical links between segments of discourse on the basis of the speaker's background semantic and pragmatic knowledge. Discourse representations produced by DICE are in the form of segmented DRSs (SDRSs) (Asher, 1993), in which discourses are represented as DRSs plus discourse relations. These discourse relations act as constraints on discourse coherence by

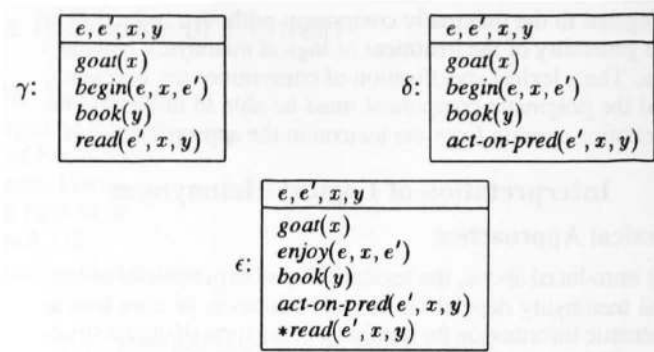


Figure 1: Logical forms associated with the three discourse continuations in (9c).

constraining the semantic content of the DRSs they connect. For example, if the discourse relation *Narration*, which conveys that one constituent of the discourse is a consequent of a previous constituent, is to hold, it must be possible to compute a common topic between the two constituents. This explains the incoherence of the discourse in (10) — there is no topic shared between the two sentences in the discourse.

- (10)*Max came in. Mary's hair is black.

For reasons of space, I will not be able to give a formal account of the analysis of the discourses in (9). Rather, I will give an informal description of the analysis. The reader is referred to Lascarides and Asher (1991, 1993), Lascarides and Copestake (1995), and Lascarides *et al* (1995) for details of the formal application of DICE.

I label the DRSs representing (9a,b) as α, β and the DRSs representing (9c(i-iii)) as γ, δ, ϵ respectively in Figure 1. The DRS γ will represent the infeasible interpretation *He began reading your book at 9pm* for (9c(i)), because the sentence has the form *begin+NP* and so the default telic event *read* is promoted to infeasible via DefFill in this case. In contrast, δ will represent the interpretation *He began doing something with your book* (there is no default interpretation for what the event done with the book is in the case of *begin on*) and ϵ represents the defeasible interpretation *He enjoyed reading your book*. The * in the DRS ϵ marks the fact that the telic event is a default specification. It has remained default after lexicosyntactic processing due to the nature of *enjoy+NP*. The process by which these interpretations are determined is described in detail in Verspoor (1996).

Let us consider how the rules apply to the discourse (9a,b,c(i)). β must be attached to α . Assuming that DICE calculates that *eating everything in his cage* is a subtype of *going nuts*, we can attach α to β with the discourse relation *Elaboration*, which conveys, as intuitions would dictate, that the event in β is a part of the event in α . Now we must incorporate γ into the SDRS *Elaboration*(α, β). We can attempt to attach γ to either α or β with a discourse relation.

Attaching γ directly to α does not seem to be possible. Intuitively, there is no generalization between α :*going nuts* and γ :*beginning to read your book*, and so there does not seem to be a topic common to the two structures which is consistent with world knowledge about *going nuts*. Attaching γ to α with *Narration* is therefore not possible. Furthermore, the latter

event is not a subtype of the former, and so *Elaboration* is not possible either. Indeed, the rules in DICE will compute that no discourse relation can adequately relate them on the basis of semantic and pragmatic knowledge. Attaching γ to β also fails in a similar manner. Thus, there is no way to make sense of (9c(i)) given the preceding context and the indefeasible interpretation resulting from lexicosyntactic processing, as (9a,b,c(i)) is predicted in DICE to be an incoherent discourse.

We can contrast this with the discourses (9a,b,c(ii)) and (9a,b,c(iii)). In both discourses, we have the same first step as above: β is attached to α with the relation *Elaboration*. Subsequently, the DRSs for the continuation of the discourses, δ and ϵ respectively, must be attached. This is done in both cases according to the analysis in L&C (1995) (for details see that paper). Since δ provides only an underspecified interpretation, the context serves to specify *act-on-pred*(e', x, y) to *eat*(e', x, y). This occurs as a result of the constraints imposed by the rhetorical relations; the strongest coherence for the discourse results when δ is in a clear rhetorical relation to the previous discourse. Here, *doing something with your book* has no clear relation to *eating everything in his cage* whereas *eating your book* is a subtype of that event and so the *doing something with your book* event is specified to *eating your book*. Similarly, for ϵ the context overrides the default interpretation *read*(e', x, y) with *eat*(e', x, y), via the axiom *Discourse Wins* combined with the preference for a strongly coherent discourse. In this case, *eating your book* is in a much stronger rhetorical relation to *eating everything in his cage* than *reading your book* is and so the former is preferred. For each of these discourses, then, an interpretation of the continuing sentence is established which would allow DICE to attach the sentence to the discourse via a clear rhetorical relation. This explains the felicity of these discourses in contrast to the discourse (9a,b,c(i)).

Infeasibility of a purely pragmatic explanation

I have shown that lexical specification of defaults combined with syntactic control over the persistence of these defaults into the pragmatic component facilitates interpretation of logical metonymies in a discourse context. A purely pragmatic explanation (e.g. Hobbs *et al.*, 1993) of such data, i.e. an explanation not relying on lexicosyntactic factors but only on world knowledge and contextual influences, would fail to account for the incoherence of the discourses (3a,b) and (9a,b,c(i)) in contrast to the coherence of the discourses (3a,c) and (9a,b,c(ii)).

The semantic representations for *begin the book* and *begin on/with the book* in such an approach would be identical — they would both correspond to the logical form δ in Figure 1, as there would be no lexical specification of default interpretations assumed. There would therefore be no basis for distinguishing their behavior with respect to the influence of discourse, even if one wanted to assume that the ability of discourse information to influence the interpretation of a particular construction were specified pragmatically. One probably does not wish to make such an assumption in any case since the primary determining factor of this behavior seems to be syntactic. Furthermore, the specification of default interpretations of logical metonymies in the pragmatic component would result in a great loss of generality, because information such as the relationship between possible default interpreta-

tions and the semantics of the noun in the NP complement (i.e. the fact that default interpretations correspond to the telic or agentive roles) could not be captured in any straightforward way, and also because there are several different types of metonymies (not only *verb+NP* but also *adjective+noun*) which display the same interpretation patterns. A pragmatic approach would be forced to specify the default interpretation of each individual logical metonymy in an *ad hoc* manner.

Conclusions

A lexically-driven approach to logical metonymy allows predictions about the range of interpretations for these constructions and the defeasibility or indefeasibility of those interpretations to be captured in a general way. The definition of a pragmatic component which has access to this lexical information is critical to the modeling of the behavior of logical metonymies in discourse contexts. I have shown that the infelicity of certain logical metonymy constructions in some discourses depends on the non-default nature of the lexicosyntactically determined interpretation for such constructions. When a non-default interpretation for a sentence cannot be coherently tied in to the discourse in which the sentence appears, discourse information cannot override that interpretation with a more coherent one, and so the sentence is judged infelicitous in that discourse — the discourse as a whole is weak. This work emphasizes the complex nature of the interaction between lexicosyntactic and pragmatic processing; discourse-level analysis is often constrained by lexical properties of the constituents of the discourse.

An open issue which remains is why it is that verb phrases differ in the persistence of their default interpretations into the pragmatic component.

Acknowledgements

This research was supported by an award from the Engineering and Physical Sciences Research Council, and by a studentship from the World Bank. The author would also like to thank Alex Lascarides, Bernie Jones, and the CogSci'96 reviewers for valuable feedback on the work.

References

- Asher, N. (1993). *Reference to Abstract Objects in Discourse*. Kluwer Academic Pub.
- Asher, N. and Lascarides, A. (1995). Lexical disambiguation in a discourse context. *Journal of Semantics*, 12(1):69–108.
- Copestake, A. (1992). *The representation of lexical semantic information*. PhD thesis, Sussex University. Cognitive Science Research Paper CSR 280.
- Copestake, A. (1993b). The compleat LKB. ACQUILEX-II Deliverable 3.1.
- Copestake, A. and Briscoe, T. (1995). Semi-productive polysemy and sense extension. *Journal of Semantics*, 12(1):15–68.
- Godard, D. and Jayez, J. (1993). Towards a proper treatment of coercion phenomena. In *Sixth Conference of the European Chapter of the ACL*, pages 168–177. Utrecht.
- Hobbs, J. R., Stickel, M. E., Appelt, D. E., and Martin, P. (1993). Interpretation as abduction. *Artificial Intelligence*, 63:69–142.

- Lascarides, A. and Asher, N. (1991). Discourse relations and defeasible knowledge. In *Proceedings of the 29th Annual Meeting of the Association for Computational Linguistics (ACL-91)*, pages 55–63, Berkeley, California.
- Lascarides, A. and Asher, N. (1993). Temporal interpretation, discourse relations and common sense entailment. *Linguistics and Philosophy*, 16(5):437–493.
- Lascarides, A., Briscoe, E., Asher, N., and Copestake, A. (1996a). Persistent order independent typed default unification. *Linguistics and Philosophy*, 19(1).
- Lascarides, A. and Copestake, A. (1995). The pragmatics of word meaning. In *Proceedings of Semantics and Linguistic Theory V*.
- Lascarides, A., Copestake, A., and Briscoe, E. (1996b). Ambiguity and coherence. *Journal of Semantics*, 13(2).
- Pustejovsky, J. (1991). The generative lexicon. *Computational Linguistics*, 17(4).
- Pustejovsky, J. (1995). *The Generative Lexicon*. MIT Press, Cambridge, MA.
- Pustejovsky, J. and Bouillon, P. (1995). Aspectual coercion and logical polysemy. *Journal of Semantics*, 12(2):133–162.
- Verspoor, C. M. (1996). Conventionality-governed logical metonymy. Unpublished ms.