

RESOLVING A PRAGMATIC PREPOSITIONAL PHRASE ATTACHMENT AMBIGUITY

Christine H. Nakatani

Department of Computer and Information Science, University of Pennsylvania, Philadelphia, PA 19104
email: nakatani@linc.cis.upenn.edu

1. Introduction

To resolve or not to resolve, that is the structural ambiguity dilemma. The traditional wisdom is to disambiguate only when it matters in terms of the meaning of the utterance, and to do so using the computationally least costly information. NLP work on PP-attachment has followed this wisdom, and much effort has been focused on formulating structural and lexical strategies for resolving noun-phrase and verb-phrase (NP-PP vs. VP-PP) attachment ambiguity (e.g. [8, 11]). In one study, statistical analysis of the distribution of lexical items in a very large text yielded 78% correct parses while two humans achieved just 85%[5]. The close performance of machine and human led the authors to pose two issues that will be addressed in this paper: is the predictive power of distributional data due to “a complementation relation, a modification relation, or something else”, and what characterizes the attachments that escape prediction?

2. Pragmatically ambiguous PPs

Although structural and lexical rules alone do not suffice to disambiguate all kinds of PPs, discourse modelling is viewed as computationally costly (cf. [1]). The debate over resolution strategies is not simply about practicality, but rather, at stake is the notion of what exactly it means for a PP to attach. This paper defends discourse-level strategies by arguing that a certain PP-attachment ambiguity, sentential vs. verb-phrase (S-PP vs. VP-PP), reflects a third kind of relation that is pragmatic in nature. As noted in [11], context-dependent preferences cannot be computed *a priori*, so pragmatic PP-attachment ambiguities are among those that defy structural and lexical rules for disambiguation.

Another criticism aimed at discourse-level approaches is that pragmatic ambiguities can be left unresolved because they do not affect the meaning of an utterance. In the case of S-PPs and VP-PPs, however, the linguistic evidence points to significant meaning differences (section 3). This paper offers a unified account of the linguistic behavior of these PPs which is expressed in a new formalism (section 4), and concludes that the resolution of pragmatic PP-attachment ambiguity is necessary for language understanding (section 5).

3. The need to disambiguate

3.1 Linguistic evidence

Linguists have identified instrumental, locative and temporal adverbial PPs as the most structurally unrestricted, context-dependent types of PPs [6, 10]. These kinds of PPs often can attach either to S or VP. Thus, *Warren sang in the park* can be paraphrased as either *Where Warren sang was in the park* or *What Warren did in the park was sing*. Kuno argues that the former interpretation involves a place-identifying VP-PP, and the latter a scene-setting S-PP. Also, the following meaning differences occur:

given-new/theme-rheme S-PPs are given/themes, VP-PPs are new/rhemes.

preposability S-PPs can be preposed, preposed VP-PPs sound awkward and often change meaning.

entailments S-PP utterances have no entailments of the utterance without the PP. For VP-PPs, the utterance without the PP is entailed only if the utterance is affirmative.

negation S-PPs always lie outside the scope of negation, VP-PPs may or may not lie inside the scope of negation.

These aspects of meaning cannot be dismissed as spurious. Consider Kuno's pair of sentences:

- Jim didn't visit museums in Paris, but he did in London (1).
- Jim didn't visit museums in Paris: he visited museums in London (2).

Kuno assigns (1) the interpretation in which the PPs are sentential and two events are described: although Jim visited museums only in London, he also went to Paris. Sentence (2) is assigned the reading that Jim was not in Paris at all but went only to London where he visited museums. The PPs are verb-phrasal and only one event is being talked about.

3.2 A pragmatic relation

The behavior of these adverbial PPs reflects neither a complementation nor a modification relation. If attachment is dictated by complementation, an instrumental PP should always appear as an argument of the verb predicate in logical form. But this sacrifices entailments for affirmative VP-PP utterances; 'butter(toast,knife)' does not logically entail 'butter(toast)' [2, 3]. If construed as a modification relation, attachment is redundant with phrase structure information and curiously depends on whether the subject, or any other constituent outside the VP, is or is not modified by the PP. There may well be reasons to preserve these relations in the syntactic structure, but they are not the relations that describe the behavior of pragmatically ambiguous PPs.

The linguistic evidence suggests that the S-PP vs. VP-PP distinction reflects a pragmatic relation, namely a *discourse entity specification* relation where *specify* means to refer in a model [4]. Since this relation cannot be represented by traditional phrase structure trees, the meaning differences that distinguish the two kinds of PPs must be captured by a different formal structure. The proposed *event formalism* treats utterances with adverbial PPs as descriptions of events and is adapted from Davidson's logical form for action sentences [2] using restricted quantification.

4. A unified formal account

4.1 Event representations

Davidson's logical form consists of an existentially quantified event entity variable and predication, as in $(\exists e)(Agt(Jones, e) \wedge Act(butter, e) \wedge Obj(toast, e) \wedge Instr(knife, e))$ for *Jones buttered the toast with the knife*. Davidson assigns equal status to all modifiers, thereby allowing events, like objects and people, to be described by any combination of their properties. This flattening of the argument structure clears the way for using restricted quantification to 'elevate' some predicates to event-specifying status. Following [12], the structure $\exists eP$ restricts the range of e to those entities

that satisfy P , an arbitrarily complex predicate of the form $\lambda u P_1(x_1, u) \wedge \dots \wedge P_m(x_m, u)$. In expressions of the form $(\exists e : \lambda u P_1(x_1, u) \wedge \dots \wedge P_m(x_m, u))[R_1(y_1, e) \wedge \dots \wedge R_n(y_n, e)]$, event-specifying predicates appear in the λ -expression while the other predicates remain in the predication Re . Hereafter, the term *event description* refers to the λ -expression, and *event predication* to the sentence predicate Re . The two parts together comprise an *event representation*.

4.2 Applying the formalism

In the formalism, (3) represents sentence (1) and (4), (2):

$$(\exists e : \lambda u \text{Agt}(J, u) \wedge \text{Loc}(P, u)) \neg [\text{Act}(v, e) \wedge \text{Obj}(m, e)] \wedge (\exists e : \lambda u \text{Agt}(J, u) \wedge \text{Loc}(L, u)) [\text{Act}(v, e) \wedge \text{Obj}(m, e)] \quad (3)$$

$$\neg (\exists e : \lambda u \text{Agt}(J, u) \wedge \text{Act}(v, u) \wedge \text{Obj}(m, u)) [\text{Loc}(P, e)] \wedge (\exists e : \lambda u \text{Agt}(J, u) \wedge \text{Act}(v, u) \wedge \text{Obj}(m, u)) [\text{Loc}(L, e)] \quad (4)$$

In (3), the thematic S-PPs (in bold) are represented in the event descriptions, whereas in (4), the nonthematic VP-PPs are in the event predications. Now the well-worn given-new distinction can be replaced by the more precise distinction made by the event formalism. *Event-specifying PP*s appear in the event description and contribute to the specification of an event entity in the discourse model. *Predication PP*s appear in the event predication and convey new information about the specified entity.

The formalism shows how preposing a VP-PP can change the meaning of the utterance. If the PP in (2) are preposed, as in *In Paris, Jim didn't visit museums: in London, he visited museums*, the original reading is lost. This is shown in the representation: $\neg(\exists e : \lambda u \text{Agt}(J, u) \wedge \text{Act}(v, u) \wedge \text{Obj}(m, u) \wedge \text{Loc}(P, u)) \wedge (\exists e : \lambda u \text{Agt}(J, u) \wedge \text{Act}(v, u) \wedge \text{Obj}(m, u) \wedge \text{Loc}(L, u))$. Since the event descriptions conflict—one event cannot take place in two places—this sentence can no longer be understood as describing a single event.

The formalism also shows different effects of negation on event-specifying and predication PP's. Sentence (2) denies the existence of any 'Jim visiting museums in Paris' event, so the quantifier lies within the scope of negation in (4). In (3) negation scopes only the event predication; sentence (1) expresses a negative fact about one event, and an affirmative fact about another. In general, a PP that lies outside the scope of negation appears in the description Pu of a representation of form $(\exists e : \lambda u Pu) \neg [Re]$. A PP that lies inside appears in the predication Re of form $\neg(\exists e : \lambda u Pu)[Re]$.

Finally, the formalism lends insight into differences in entailments. The following entailment relationship holds for affirmative VP-PP sentences, where $R_n(y_n, e)$ represents the PP predicate: $(\exists e : \lambda u Pu)[R_1(y_1, e) \wedge \dots \wedge R_{n-1}(y_{n-1}, e) \wedge R_n(y_n, e)] \Rightarrow (\exists e : \lambda u Pu)[R_1(y_1, e) \wedge \dots \wedge R_{n-1}(y_{n-1}, e)]$. A PP predicate $R_n(y_n, e)$ in a negated event predication may or may not be negated, so the entailment for negative VP-PP sentences is blocked: $(\exists e : \lambda u Pu) \neg [R_1(y_1, e) \wedge \dots \wedge R_{n-1}(y_{n-1}, e) \wedge R_n(y_n, e)] \not\Rightarrow (\exists e : \lambda u Pu) \neg [R_1(y_1, e) \wedge \dots \wedge R_{n-1}(y_{n-1}, e)]$. Why S-PP sentences have no entailments is a separate matter. Eliminating an event-specifying PP from an event description yields a representation with a different description. Intuitively, it seems desirable that no entailment relations hold between different types of entities. The formalism preserves this condition.

The proposed formalism succeeds in capturing the discourse entity specification relation and lends itself naturally to processing in an NLP system that takes seriously the dynamic nature of context. Such a system would for each utterance construct an event representation, search for a discourse entity that satisfies the event description, and use the event predication to update the information about that entity in the discourse model.

5. Conclusion

A preliminary algorithm for processing highly ambiguous PP's has been worked out in [7]. The algorithm uses intonation [9], centering and word order information to construct and process event representations in a discourse model structured after [4]. The wider applicability of the two-part event formalism has not yet been tested. Nevertheless, one conclusion is that the value of resolving any structural ambiguity can only be measured in terms of the semantics of the structural formalism itself. In the case of VP-PP vs. S-PP ambiguity, an NLP system must not idly wait for syntax to choose how a PP should pragmatically function. The traditional wisdom—find the meaning and do so efficiently—instead suggests that more productive than demanding of syntax unreasonably diverse expressive powers is to search for direct linguistic correlates of pragmatic meaning that can be efficiently encoded in a dynamic pragmatic formalism.

Acknowledgements

The author thanks Barbara Grosz and Julia Hirschberg, who both advised this research, for valuable comments and guidance; and acknowledges current support from a National Science Foundation Graduate Fellowship. This paper stems from research carried out at Harvard University and at AT&T Bell Laboratories.

References

- [1] Altmann, G. and M. Steedman 1988. Interaction with context during human sentence processing, *Cognition*, 30(3).
- [2] Davidson, D. 1967. The logical form of action sentences, in Davidson and Harman, eds., *The Logic of Grammar*, pp. 235-246, Dickenson Publishing Co., Inc., Encino, CA, 1975.
- [3] Fodor, J. A. 1972. Troubles about actions, in Harman and Davidson, eds., *Semantics of Natural Language*, pp. 48-69, D. Reidel, Dordrecht-Holland.
- [4] Grosz, B. J. and C. Sidner 1986. Attention, intentions, and the structure of discourse, *CL*, 12(3).
- [5] Hindle, D. and M. Rooth 1990. Structural ambiguity and lexical relations, *Proceedings of the DARPA Speech and Natural Language Workshop*, Hidden Valley, Pennsylvania.
- [6] Kuno, S. 1975. Conditions for verb phrase deletion, *Foundations of Language*, 13.
- [7] Nakatani, C. 1990. A discourse modelling approach to the resolution of ambiguous prepositional phrases, manuscript.
- [8] Pereira, F. C. N. 1985. A new characterization of attachment preferences, in Dowty, Karttunen and Zwicky, eds., *Natural Language Parsing*, pp. 307-319, Cambridge University Press, Cambridge.
- [9] Pierrehumbert, J. and J. Hirschberg 1990. The meaning of intonational contours in the interpretation of discourse, in Cohen, Morgan and Pollack, eds., *Intentions in Communication*, pp. 271-311, MIT Press.
- [10] Reinhart, T. 1983. *Anaphora and Semantic Interpretation*, University of Chicago, Chicago.
- [11] Shieber, S. 1983. Sentence disambiguation by a shift-reduce parsing technique, *Proceedings of 21st Meeting of the ACL*, Cambridge, MA.
- [12] Webber, B. 1983. So what can we talk about now?, in Brady and Berwick, eds., *Computational Models of Discourse*, pp. 331-371, Cambridge, MA, MIT Press.