Some Pragmatic Issues in the Planning of Definite and Indefinite Noun Phrases

Douglas E. Appelt

Artificial Intelligence Center, SRI International
and
Center for the Study of Language and Information
Stanford University

1 Introduction

In this paper we examine the pragmatic knowledge an utterance-planning system must have in order to produce certain kinds of definite and indefinite noun phrases. An *utterance-planning system*, like other planning systems, plans actions to satisfy an agent's goals, but allows some of the actions to consist of the utterance of sentences. This approach to language generation emphasizes the view of language as action, and hence assigns a critical role to pragmatics.

The noun phrases under consideration in this paper are those that presuppose the existence of an individual that could be described by the description D . In other words, when a speaker uses a noun phrase with description D , it makes sense to ask the question "Which x is D ?" This criterion includes more than strictly referential uses of noun phrases, because it is not necessary for the speaker or hearer to *know* what individual is described by D — it is merely necessary that the existence of such an individual is presupposed. Consider the attributive description in sentence (1):

- (1) The runner who wins tomorrow's race will qualify for the semifinals.

The description "runner who wins tomorrow's race" cannot be referential, because, under ordinary circumstances, the speaker could not possibly know who it is that would fit the description. Nevertheless, it is still reasonable to ask which runner will win tomorrow's race, because the description is objectively true of some individual.

This qualification excludes noun phrases whose referents are bound within the scope of a universal quantifier, such as "the woman ..." in

- (2) Every man wants to meet the woman of his dreams.

For a similar reason, indefinites within the scope of a sentential negation are excluded because they introduce an existential quantifier, which, under the scope of negation,

is really a universal quantifier. Therefore, "a screwdriver" in

- (3) John does not have a screwdriver.

is excluded because, under most circumstances of its use, there is no screwdriver that the description in sentence (3) denotes. *Predicate nominals* are excluded, as in the sentence

- (4) John wants to be a doctor.

because one would not ask the question "Which doctor does John want to be?"

The choice of this particular class of noun phrases is motivated by considerations relevant to planning. When a speaker communicates with a hearer, he often intends the hearer to hold some attitudes toward individuals in the domain. This is particularly true in task-oriented dialogues where the hearer may have to locate and manipulate things in his environment.

The theory of utterance planning assumed for the purpose of this analysis is the one embodied in KAMP (Appelt, 1985). Individuals are represented by terms in an intensional logic of knowledge and action. A metalanguage is used to axiomatize the relationship that holds between the terms and the individuals they denote. The terms can consist of predicates combined with an iota operator, as in $\iota x D(x)$, where

$$D(x) \equiv D_1(x) \wedge \dots \wedge D_n(x).$$

The predicates D_i are called *descriptors*, and their conjunction, D , is called a *description*. Because most noun phrases employ terms that are constructed from a description, often the words "term" and "description" are used interchangeably.

The propositional content of the speaker's utterance is represented by a sentence in the intensional logic involving the terms discussed above. Uttering a sentence entails performing a number of actions, called *concept activation actions*, which result in the terms constituting the proposition receiving a special status called "*active*." The proposition

that the speaker intends to convey is a predication involving the active terms. Referring is a particular type of concept activation action with relatively strict conditions on what must be mutually believed by the speaker and hearer for the action to succeed. Searle (1969) presents an analysis of referring as a speech act and dismisses many uses of noun phrases as nonreferring. Such nonreferring noun phrases occur very frequently, and the considerations that underly their planning share much in common with those that underly actual referring. Therefore, the concept activation action provides a suitable generalization that allows a plan-based treatment of many more uses of noun phrases.

2 Research Objectives

The analysis presented in this paper represents one of the first steps toward a plan-based account of definite and indefinite noun phrases. Ideally, such an account would (1) provide a semantics for noun phrases, (2) define an actions like "uttering a definite noun phrase," and (3) provide an analysis that shows how the speaker's intentions follow directly from the semantics of the noun phrase he utters, plus conditions on mutual knowledge and general principles of rationality. This program is very much in the spirit of the analysis of illocutionary acts provided by Cohen and Levesque (1980), who demonstrate how illocutionary acts can be defined in terms of the kinds of inferences made, given a semantic analysis of an utterance, facts about mutual knowledge, and general principles of rational behavior.

Cohen (1984) provided such an analysis for referring actions by postulating a semantics for the definite determiner that would give the semantics of a definite noun phrase as a request to identify the referent of a description. This analysis would be impossible to extend to the more general concept activation actions, because, in some cases, the speaker intends that the hearer *not* identify the denotation of the description, even when a definite noun phrase is used. A complete analysis along these lines that subsumes both referring and nonreferring noun phrases has yet to be worked out.

As an intermediate step toward this ultimate goal, we shall propose a taxonomy of concept activation actions that convey the various intentions a speaker may have with respect to a hearer and a description. This taxonomy is of theoretical interest, because it characterizes differences and similarities among uses of noun phrases that current theories do not characterize. It is also of practical interest for utterance planning, because the set of actions to be proposed provides a useful level of abstraction for the reasoning processes of an utterance-planning system. For example, certain planning strategies such as action subsumption (Appelt, 1985) are applicable only to certain kinds of concept activation actions and not to others. Therefore, even if the complete plan-based analysis of noun phrases is worked out, the taxonomy of actions presented here will still be of

practical importance.

Until an analysis like Cohen and Levesque's is worked out, the concept activation actions here will be treated like illocutionary acts in a speech-act theory. When a hearer understands an utterance, he reasons about whether it constitutes an assertion, a request, a warning, etc. Therefore, understanding one of the definite or indefinite noun phrases under consideration in this paper is assumed to entail recognition of what concept activation action the speaker intends to perform.

3 Summary of Actions Underlying Noun Phrases

There are many distinctions that one could draw between noun phrases, only some of which are relevant to planning. For example, one could distinguish noun phrases that refer to amorphous substances from those that refer to discrete entities. Such a distinction may have some valid motivation, but it is not necessarily so from the standpoint of planning. It would be well motivated only if there were a clear difference in the preconditions and effects of the concept activation actions underlying mass terms, or in the strategy for the selection of descriptors. This does not seem to be the case for mass versus discrete entities.

However, there are two criteria that clearly affect the relevant preconditions, intended effects, and planning strategies of concept activation actions: (1) whether the speaker intends that the hearer *identify* the denotation of the description, and (2) how much mutual knowledge the speaker and hearer share about the description's denotation. The first criterion is what (roughly) distinguishes referring noun phrases from nonreferring noun phrases. The necessity of the hearer performing the identification constrains the description to be one that facilitates the hearer's formulation of a plan to do so.

The second criterion is the knowledge that is shared by the speaker and the hearer at the time of the utterance. Planning strategies are influenced by whether or not the speaker and hearer mutually believe appropriate facts about the intended referent. In particular, if the speaker and hearer share enough knowledge about the description's denotation and the contextual situation, it may be possible for the hearer to recognize the speaker's intentions using only a subset of the descriptors in the noun phrase's description. In such a situation, the speaker can augment the description with additional descriptors for the purpose of informing the hearer that they are true of the denotation of the other part of the description. Such a strategy is called *action subsumption*, (Appelt, 1985). The action subsumption strategy cannot be used with concept activation actions that are not based on shared knowledge.

Since there are two dimensions relevant to characterizing concept activation actions, it is possible to define four

Figure 1: Four Types of Concept Activation Actions

actions, as illustrated in Figure 1. These actions are SI (shared concept activation with identification intention), NSI (nonshared concept activation with identification intention), SNI (shared concept activation with no identification intention), and NSNI (nonshared concept activation with no identification intention.) Each action has distinct preconditions, effects, and an associated planning strategy.

4 Mutual Knowledge and Identification

The two most important considerations in planning concept activation actions are (1) whether or not the speaker intends the hearer to *identify* the referent of the description and (2) what *knowledge* about the description's possible denotations the speaker and hearer share.

What it means for an agent to "identify" the referent of a description is a topic of considerable complexity. Searle (1969) states that "So identification ... rests squarely on the speaker's ability to supply an expression ... which is satisfied uniquely by the object to which he intends to refer." What counts as an identifying description depends on the purpose for which the agent is identifying the description's denotation. For example, the description that one must know to carry out a plan requiring the identification of "John's residence" may be quite different depending on whether one is going to visit him, or mail him a letter. If I want to speak to a guest at a Halloween party, I need only a description capable of distinguishing him from the other guests at the party, not to know who it really is wearing the disguise.

Identification of the denotation of a term *D* is therefore defined as finding another term *D'* (called a *prima facie* (PF) identifiable term) that has the same denotation as *D* according to the hearer's knowledge, but that meets certain

syntactic criteria for being the "right kind" of term. It is stipulated that any two distinct PF identifiable terms must denote different individuals in the same situation. The simplest criterion for PF identifiability that meets this requirement is that the term be a *standard name*. Because each standard name denotes the same individual in any context, knowing that a particular standard name is equivalent to a term implies that the agent knows the denotation of the term. Furthermore, any two distinct standard names denote different individuals.

The standard name approach was taken by the KAMP system. The standard name assumption has two difficulties. First, it is extremely implausible to believe that an agent has a unique name for anything that can be referred to. Also, knowing a standard name implies having made an absolute identification. Therefore, to refer to a guest at a costume party, it is a consequence of successful identification that the speaker and the hearer mutually know the identity of the person in the disguise, which is obviously too strong a condition for successful reference. Developing adequate criteria for PF identifiable terms is an important research problem; however, none of the points in this paper depend on what the criteria for PF identifiability are.

The importance of mutual belief to the successful use of referring expressions was demonstrated by Clark and Marshall (1981). It was shown by a series of rather complex examples that, if one did not observe an infinite number of preconditions of the form "A believes that B believes that A believes that B believes ... description *D* applies to *R*," then it is impossible to guarantee that description *D* can be used to refer felicitously to *R*, because it would always be possible to construct some set of circumstances in which the hearer would believe the speaker intended to refer to something else. Perrault and Cohen (1981) show that a slightly weaker condition is adequate: the mutual belief preconditions have to hold in all but a finite number of cases. Nadathur and Joshi (1983) adopt a strategy that amounts to assuming that if *D* is believed to apply to *R*, then it is also mutually believed to apply to *R* unless there is reason to believe that it is not.

The case for some form of mutual belief as a prerequisite to a successful referring action is strong; however, speakers often use noun phrases that should be analyzed as referential in which it is clear from the context that, not only is the description not mutually believed to hold of the intended referent, but the speaker knows this is the case when he plans the utterance. For example, consider a situation in which the speaker is giving instructions to the hearer and says

- (5) Turn left at *the third block past the stoplight*.

This utterance might be reasonable even if the hearer had never been to the intersection in question and the speaker and hearer have no mutual belief at the time of the utterance about the location to which the speaker intends to

refer. The hearer knows that the speaker can formulate a plan at the time of the utterance that will guarantee that he will have identified the referent of the description at the time that it is needed.

This observation is one motivation for the distinction drawn along the horizontal axis of Figure 1. There are really two kinds of definite referring actions: one is that in which the precondition is mutual knowledge of a description, and the other in which there is mutual knowledge of a plan incorporating the description to acquire additional knowledge.

5 Definitions of Concept Activation Actions

This section discusses each of the four major types of concept activation actions outlined in Section 3. The definitions of the actions are not stated rigorously, but are intended to give the reader an intuitive understanding of their preconditions and effects, and how they differ from each other.

5.1 Shared Concept Activation with Identification Intention (SI)

These actions are the only type of concept activation actions that were considered in the earlier KAMP research. SI actions are used most frequently in referring to past events and objects that are not perceptually accessible to the hearer. In such situations, the hearer can perform few, if any, actions to acquire more knowledge that would enable him to identify the referent of a description whose referent was not already mutually known at the time of the utterance.

SI Action: The speaker S performs action SI with hearer H and term D

Preconditions: There is some term D' which is PF identifiable; S and H mutually believe that $\text{Denotation}(D) = \text{Denotation}(D')$.

Effect: H knows that S intends that the term D' be active.

The preconditions of this action depend strictly on the mutual belief of the speaker and the hearer at the time of the utterance. The noun phrase in a sentence such as

(6) Use the same wrench you used to unfasten the pump.

must arise from this type of action in normal situations of its use, because the description, based on a past event, does not facilitate any kind of plan for acquiring more information.

When planning an utterance, the speaker knows the PF identifiable term, and his problem is to get the hearer to

recognize the same term. Consistency with the Gricean maxim of quantity requires that the planned description be as simple or efficient as possible. There are several ways to measure the complexity of a description, including the number of descriptors involved and the ease with which these descriptors can be incorporated into the utterance. When planning an SI action, the planner's most important task is reasoning about the *efficiency* of the description.

Concept activation actions that involve shared belief about the denotation of the description at the time of the utterance have the property that they are candidates for *action subsumption*. Because the information required to perform the identification can be communicated through a subset of the descriptors in the noun phrase, or extralinguistically through pointing actions or strong contextual cues, and because the precondition

$$\text{Denotation}(D) = \text{Denotation}(D')$$

is known to hold, the speaker can use the additional descriptors to *inform* the hearer that the descriptors are true of the intended referent.

5.2 Nonshared Concept Activation with Identification Intention (NSI)

This action is what a speaker does when he wants to refer to an object that is not known to the hearer, or for which the speaker and hearer do not mutually believe enough properties at the time of the utterance so that identification can take place based on mutual knowledge.

NSI Action: The speaker S performs action NSI with hearer H and term D .

Preconditions: S and H mutually believe that there is some plan P such that, if H executes P , then in the resulting state, there exists a PF identifiable term D' such that H knows that $\text{Denotation}(D) = \text{Denotation}(D')$, and S intends that H execute P .

Effects: H knows that S intends that D be active.

The NSI action is used in situations in which the speaker and hearer do not mutually know the denotation of the description, yet, to realize the perlocutionary effects of the utterance, the hearer must be able to identify the speaker's intended referent. This lack of mutual knowledge may occur if the speaker can identify the referent from the description, but the hearer cannot, as is most likely the case in example (5). Also, as is the case in example (7), the speaker may not be able to identify the referent, but nevertheless knows of a plan the hearer can execute that will lead to the identification of the referent at the appropriate time.

(7) Get me the largest tomato from the garden.

The speaker of sentence (7) is uttering an attributive description, because he is probably not referring to a particular tomato, but to whatever tomato fits his description. However, it is conceivable that he had a particular tomato in mind, and chose that description because he believed it would lead to the best plan for the hearer to identify it, and would, in that case, be referential. One can see from this example that the referential-attributive distinction is orthogonal to the distinctions motivated by utterance planning. In both referential and attributive cases, the speaker knows that the right conditions on mutual knowledge are not satisfied for an SI action, and plans a description that he knows the hearer can use successfully. It does not matter to the planner whether the description is referential or attributive — the same reasoning takes place in both cases with the same results.

The NSI action depends on the hearer's ability to find the plan P . Therefore, the speaker must plan to furnish information as part of D that will make it as easy as possible for the hearer to formulate his plan. If the hearer has enough information to formulate P , then D is a *useful* description. It is possible for a speaker to formulate a description that, although it denotes the individual the speaker has in mind, is not useful because there is no plan the hearer can formulate to take advantage of the description. An example of such a nonuseful description would be if S and H are riding a bus, H asks at what stop he should get off, to which S replies "one stop before I do." The description "one stop before I do," while being true of a unique location, is not a useful description, assuming that the hearer has recourse only to observing the speaker's actions.

The reader may wonder if an SI action can be regarded as a degenerate case of the NSI action. In the case of the NSI action, the speaker and hearer mutually know of a plan that will result in identification of the intended referent, and in the case of the SI action, the plan is simply to do nothing, because the referent of the term is already mutually known. This is not the case, because the precondition of the SI action is that the speaker and hearer mutually believe both the description in the noun phrase and the PF identifiable description. In the case of the NSI action, the speaker and hearer mutually believe that executing plan P will result in the hearer acquiring the required information, but, since only the hearer is actually executing the plan, the speaker and hearer may never meet the mutual belief condition of the SI action. Therefore it is possible to have an NSI action with a null plan, which is not equivalent to an SI action with the same description. For example, suppose a speaker wants a son to deliver an envelope to his father, and makes the request

(8) Give this to *your father*.

although the speaker does not know who the son's father is. In sentence (8) the speaker is using the description attributively because he has no particular individual in mind,

just whoever fits the description. Furthermore, the speaker assumes that the son is capable of identifying his own father on the basis of knowledge he already has; therefore the plan for the hearer to identify the description is to do nothing. This is different from the SI action, in which there is some individual who is mutually believed to be the hearer's father.

5.3 Shared Concept Activation with No Identification Intention (SNI)

When a speaker performs an SNI action, he provides a description, but he does not intend that the hearer try to identify its denotation. Therefore, the SNI action is not a referring action, because identification is an essential part of referring. The SNI action is used when a speaker has a belief involving some individual for whom he has a description, but not a PF identifiable description, and intends that the hearer hold the same belief.

SNI Action: The speaker S performs action SNI with hearer H and term D .

Preconditions: S and H mutually believe that there exists an individual R such that $\text{Denotation}(D) = R$.

Effects: H knows that S intends that D be active.

The primary effect of the SNI action is the same as the NSI action: it activates the term corresponding to the description D . However, because the preconditions are different, no intention to identify the description is communicated, and the ultimate effect of the action on the hearer's beliefs and intentions is therefore quite different. This type of action underlies the use of an attributive description when no identification is intended. This type of action has been discussed in the literature (Donellan, 1966; Kripke, 1977) with the situation of two people discovering Smith's badly mutilated body, and one saying "The man who murdered Smith is insane." In this situation, the speaker is informing the hearer of something about the referent of the description "man who murdered Smith," but does not know who this individual is, nor does he intend that the hearer identify him. However, there are conditions on the mutual belief of the speaker and hearer for the utterance to make sense. The speaker and hearer must mutually believe that the dead man is Smith, that he was in fact murdered, and that it was a man who killed him.

5.4 Nonshared Concept Activation with No Identification Intention (NSNI).

NSNI Action: The speaker S performs action NSNI with hearer H and term D .

Preconditions: No mutual belief preconditions.

Effects: *H* knows that *S* intends that the term *D* be active.

Unlike the SNI action, the NSNI action does not require that the speaker and hearer share any knowledge about the denotation of the description prior to the utterance. This action is used by a speaker to introduce a new individual to the discourse, without intending that the hearer associate that individual with any that he already knows about. For example, a speaker says, "I met *an old friend from high school* yesterday." The speaker does not assume that the hearer shares any knowledge of his old high school friends, nor does he intend the hearer to identify the person he is talking about. The most important consideration for the planner in this case is to include enough information in the description *D* to serve the speaker's purpose in the rest of the discourse.

NSNI actions are most frequently realized by referential indefinite noun phrases (Fodor and Sag, 1982). Such a noun phrase is indefinite, but it is clear from the context that there is some particular individual that is denoted by the description.

6 Summary

This paper has examined a class of actions called concept activation actions, in which a speaker communicates the intent that the hearer recognize a particular description. The performance of one of these actions consists of uttering a noun phrase, either in isolation, or as part of a sentence. Therefore, the noun phrases resulting from the performance of a concept activation action are, in some sense, referential, even though neither the speaker nor the hearer may know the noun phrase's denotation, either at the time of the utterance or subsequently.

While the four actions discussed in this paper account for a very important class of noun phrases, the class by no means exhausts all possibilities, and further research is needed to understand the pragmatic considerations relevant to other noun phrases. Some other noun-phrase examples were discussed earlier, including quantificational noun phrases and predicate nominals. Generics, and bare plurals will require additional analysis. There is also an extremely important class of concept activation actions that has not been discussed here, namely *coreferring actions*, which entail the activation of terms that have already been introduced to the discourse.

This analysis of the actions underlying the production of noun phrases is of particular importance to utterance planning. Planning requires a characterization of actions that describes what their effects are, when they are applicable, and what strategies are available for their expansion. The four actions described in this paper fill an important gap that has been left open in previous utterance-planning research.

Acknowledgements

This research was supported, in part, by the National Science Foundation under grant DCR-8407238 and was made possible, in part, by a gift from the System Development Foundation to SRI International as part of a coordinated research program with the Center for the Study of Language and Information at Stanford University. The author is grateful to Barbara Grosz and Ray Perrault for comments on earlier drafts of this paper.

References

- Appelt, D. E., *Planning English Sentences*, Cambridge University Press, Cambridge, UK (1985).
- Clark, H. and C. Marshall, "Definite Reference and Mutual Knowledge," in Joshi, Sag, and Webber (eds.) *Elements of Discourse Understanding*, Cambridge University Press, Cambridge, UK (1981) pp. 10-63.
- Cohen, P. R., and C. R. Perrault, "Elements of a Plan Based Theory of Speech Acts," *Cognitive Science* 3 (1979) pp. 177-212.
- Cohen, P. R., "Referring as Requesting," *Proceedings of the Tenth International Conference on Computational Linguistics* (1984) pp. 307-221.
- Cohen, P. R. and H. Levesque, "Speech Acts and the Recognition of Shared Plans," *Proceedings of the Third Biennial Conference*, Canadian Society for Computational Studies of Intelligence (1980).
- Cohen, P. R., "Pragmatics, Speaker-Reference, and the Modality of Communication," *Computational Linguistics* 10 (1984) pp. 97-146.
- Donellan, K., "Reference and Definite Descriptions," *Philosophical Review*, 75 (1966) pp. 281-304.
- Fodor, J. and I. Sag, "Referential and Quantificational Indefinites," *Linguistics and Philosophy* 5 (1982) pp. 354-398.
- Kripke, S., "Speaker Reference and Semantic Reference," in French, Uehling, and Wettstein (eds.) *Contemporary Perspectives in the Philosophy of Language*, University of Minnesota Press, Minneapolis, MN, (1977) pp. 6-27.
- Nadathur, G. and A. Joshi, "Mutual Beliefs in Conversational Systems: Their Role in Referring Expressions," *Proceedings of the Eighth International Joint Conference on Artificial Intelligence*, (1983) pp. 603-605.
- Perrault, C. R., and P. R. Cohen, "It's for Your Own Good: A Note on Inaccurate Reference," in Joshi, Sag, and Webber (eds.) *Elements of Discourse Understanding*, Cambridge University Press, Cambridge, UK (1981).
- Searle, J. R., *Speech Acts*, Cambridge University Press, Cambridge, UK (1969).