

Some Psycholinguistic Constraints on the Construction and Interpretation of Definite Descriptions¹

Andrew Ortony
Center for the Study of Reading
University of Illinois at Urbana-Champaign

Introduction

It is a curious and important fact about natural languages that they can be and often are used nonliterally. Whereas in artificial languages it is normally impossible to distinguish between the meaning of an expression on the one hand, and the intended meaning the user wishes to convey with that expression on the other (the two are identical), in natural languages a parallel distinction between sentence meaning and speaker meaning (see, for example, Searle, 1975, in press) underlies all nonliteral uses. One of the central concepts in the analysis of nonliteral uses of language is that of indirectness. It is a notion that has begun to attract the attention of linguists, philosophers, psychologists, and computer scientists in their various efforts to come to a better understanding of natural languages and of human linguistic performance. One of my purposes in this paper is to show how central a concept indirectness is with respect to the production and comprehension of definite descriptions.

The main problem with which I am concerned is a multi-level one. At the most general level it concerns the way in which people determine the referents of definite descriptions, and how language users choose the definite descriptions they do. More specifically, I am concerned with the question of the constraints that exist upon how a thing can be referred to. What makes this an interesting problem is the fact that it seems not to be necessary for a referring expression to be based on either information that has already been made explicit in the preceding discourse, nor even on information that is entailed by what has. Yet clearly, there are constraints on the expressions that can be used if there is to be a realistic hope of communicative success.

The question of what is to count as a reasonable way of referring to something in part depends for its answer on what counts as a reasonable indirect use of language. When, for example, one refers to the 1977/78 Seattle basketball team as The Cinderella of the NBA one is using a definite description based on a predicate that is not literally true of the intended referent but that is metaphorically applicable to it. As one thinks about the processes that might be involved in the production or comprehension of such an expression they appear to be very

complex, yet however complex they may be, people usually engage in them without any apparent difficulty. At present there appears to exist no adequate theoretical account of what these processes are like, perhaps because a comprehensive treatment of definite descriptions has as a prerequisite a theory of indirectness, and that in turn seems to hinge on a more comprehensive theory of speech acts than is currently available. My own proposals are not intended to fill all these gaps, but they are intended to sketch a possible direction for doing so. The main goal that I have is to suggest a way of imposing limits on indirectness, and then, to show how those same limits are needed to account for some important constraints on successful definite descriptions.

Definite Descriptions and their Textual Relations

I shall take it as axiomatic that every definite description is based upon a predicate that is supposed to be uniquely applicable (at least within the context of the discourse) to some entity relevant to the discourse. Thus, the definite description The first man on the moon is based on the predicate is/was the first man on the moon, and it is assumed to be applicable to some entity (e.g. Neil Armstrong) relevant to the discourse. It is important to note at the outset that coreferential expressions cannot always be substituted for one another without a change of meaning. For instance, if, on arriving in a strange unfamiliar hotel in a foreign land one were to utter (1a), it hardly makes sense to say that it is equivalent to uttering (1b).

- (1a) I feel like the first man on the moon.
(1b) I feel like Neil Armstrong.

However, if the speaker can safely assume that his audience knows that Neil Armstrong was the first man on the moon, (1b) could be used as an indirect way of achieving the communicative intent of (1a). Notice, it is not necessary to know who the first man on the moon was in order to fully understand (1a), whereas it is necessary to know that Neil Armstrong was the first man on the moon in order to properly understand (1b) (although one might not understand why the speaker used (1b) with its unnecessary demands on additional knowledge and inferences in preference to (1a).) In any event, it seems that even the relationship between definite

descriptions and proper names may sometimes depend on a notion of indirectness (see Ortony & Anderson, 1977).

The interpretation of definite descriptions often relies heavily on the establishment of inferential relationships of various kinds in order to determine which entity is being referred to. Such inferences tend to be forced jointly by the desire of the hearer or reader (hereafter referred to simply as "the hearer") to make sense of the discourse, and the assumption that the speaker or writer (hereafter, simply, "the speaker") is communicating in accordance with the cooperative principle (see Grice, 1975). This latter assumption is critically important in cases where the predicate underlying the definite description is not obviously true of the intended referent--and since these cases appear, at least on the surface, to constitute the most difficult ones, I shall concentrate on their analysis, to some extent at the expense of simpler examples. I shall call such cases "definite descriptions of inference." The overriding logic of the determination of the referents of such definite descriptions is that if the speaker is communicating in accordance with the cooperative principle certain assumptions have to be recognized in order for the expression in question to successfully identify the intended referent. These assumptions often serve to "sneak in" new information about the referent (in much the same way as appositive relative clauses introduce new information). The identification of these assumptions is based on inferences of various kinds.

Definite descriptions of inference can be contrasted with definite descriptions based on entailment relations. Definite descriptions based on entailment are those for which the referent can be determined either by transforming a predicate that has already appeared in the discourse into a definite description (descriptions based on the principle of identity), or by relating the description to predicates that have appeared earlier, on the basis of rules of formal logic (e.g. modus ponens) applied to them.² The important difference between a definite description of entailment and a definite description of inference is that the interpretation of the former does not depend on the provision of suppressed premises drawn from the comprehender's general world knowledge. In a definite description of inference it does. An example of a definite description of entailment can be found in (2), where the underlined expression is entailed by the content.

- (2) A well-dressed man entered the room and greeted the hostess. Although everyone else was drinking sherry, he asked the waiter for a scotch. The waiter gave him one. The man with the scotch walked over to his host.

In this case, if the waiter gave the man a scotch, it entails that the man had the scotch, and so, within the constraints imposed by the context, he can be uniquely identified by the definite description the man with the scotch.

76
With definite descriptions of inference, as with communication in general, success often depends on the speaker and the hearer sharing a common background of knowledge (see, for example, Stalnaker, 1974). Definite descriptions of inference are more complex. For example, suppose that in (2) the sentence The waiter gave him one is omitted. Then, the definite noun phrase The man with the scotch only succeeds in referring to the right man if it is assumed that the man who asked for a scotch was given one. Unfortunately, only in biblical circles is it true that asking for something guarantees being given that thing. So, in order for the hearer to identify the intended referent he has to assume that the man got his scotch. Of course, this assumption comes easily for it can be made on the basis of a plausible inference requiring only the introduction of plausible suppressed premises, such as that when a guest asks a waiter for a particular kind of drink at a cocktail party, the waiter normally obliges if that drink is available. This constitutes a simple example of a definite description of inference.

On encountering a definite description, the hearer has to assume that the description does indeed refer to some already mentioned person or thing, in this case, say, the guest. In doing so, he makes inferences that fill in what went before--that is he makes inferences about what might have been asserted to enable the predicate underlying the description to be both applicable and relevant. The comprehender might reason as follows in the present example: "[f this expression refers to the guest, then it must be the case that the waiter gave him a scotch. This is quite plausible since it is customary for waiters at cocktail parties to give guests the drinks they request if those drinks are available. It is plausible that scotch was available, since it is a frequently served drink at such occasions. So I shall assume that this is what happened and that is why the guest was referred to as 'the man with the scotch' " Whether or not people normally construct such chains of reasoning in order to identify the antecedents of definite descriptions is not the issue here. What is the issue, as we shall see later, is that it be possible to construct such a chain. Certainly, one has to suppose that the kind of general world knowledge required to do so is normally available during the comprehension process. The relevant frames, scripts, schemata, or whatever other knowledge structures are supposed, are presumably activated.

Definite descriptions of inference involving indirectness, like (3) and (4) below, tend to be more complex. They are characterized by the fact that the applicability of the predicates underlying them often depends on the utilization of knowledge that Morgan (1978) calls knowledge about the language, as opposed to knowledge of the language. These turn out to be cases of inferences involving knowledge about illocutionary forces and perlocutionary effects (see Austin, 1964).

- (3) The hostess offered the guest some cake. He told her that he was on a diet. His brother, who was with him, told her that he personally was not

on a diet. The man who had refused the cake walked over to his host.

- (4) The hostess asked the man where his wife was. He replied "Mind your own business, you old bag." The hostess was furious that the man who had insulted her had been invited to her party.

The interpretation of (3) requires not only semantic and general world knowledge in the way that (2) does, it also requires the knowledge that saying that one is on a diet can count as rejecting an offer to eat something. In the case of (4) it requires the knowledge that the violation of certain language-use conventions can count as offensive behavior. Of course, in a sense, this kind of knowledge about the conventions of language use and the social/communicative consequences of their violation is knowledge of the world, just as knowing that waiters normally serve the drinks they are asked to is. But, insofar as it is knowledge of conventions about the use of language, and insofar as this is an area which has been singled out as being of core concern in pragmatics, it is worth separating such cases from the other kinds of cases, like (2). In fact, I think, the distinction is difficult to uphold because the mechanism required to deal with indirectness is the same kind of inferential mechanism as is required to deal with "ordinary" knowledge of the world.

The question that eventually has to be answered concerns the constraints that there are on the predicates employed in definite descriptions. My view is that the answer to this question depends on finding an answer to a more general question about the pragmatics of language, namely the question: what constraints are there on what is relevant (in the sense used by Grice, 1975 and others). Staying, for the moment, with definite descriptions, compare (5) and (6) below

- (5) The hostess offered the guest some cake. He told her that he was on a diet. His brother, who was with him, told her that he personally was not on a diet. The man who thought he ought not to eat fattening things walked over to his host.
- (6)* The hostess offered the guest some cake. He told her that he was on a diet. His brother, who was with him, told her that he personally was not on a diet. The man who was not hungry walked over to his host.

It seems to me that whereas (5) is perfectly coherent, (6) is not. It becomes coherent, however, if the context is changed so that instead of (he) told her that he personally was on a diet it reads (he) told her that he personally had just eaten, then both (3) and (6) seem perfectly acceptable. It seems, then, that the appropriateness of the definite description depends on the appropriateness of its underlying predicates. Telling someone that one has just eaten is an appropriate, relevant, piece of information for

permitting the inference that one is not hungry, and/or that one does not want the offered food. By contrast, telling someone that one is on a diet is an appropriate piece of information for permitting the inference that one does not want what is being offered, but it is not appropriate for the inference that one is not hungry.

Indirect Speech Acts

A major part of my thesis is that the predicate underlying a definite description of inference is constrained by the relevance relation in just the same way as that relation constrains what counts as an indirect speech act in a discourse. If this is so, then it will help to have a working hypothesis about the constraints that exist on indirect speech acts.

Suppose the situation is that described in (3), namely, one in which someone is offered some cake and in uttering (7) intends to refuse the cake.

- (7) I am on a diet.

The question we have to answer is this. Since it does not follow logically from (7) that the intention was to refuse the cake, on what basis does a listener come to the conclusion that indeed that was the intention? Furthermore, why does, for example, (8) not succeed in communicating the refusal?

- (8) My mother is an opera singer.

Perhaps one should reject (8) on some very general grounds. For example, on the grounds that one cannot randomly assign a sentence to an intention and expect to be understood. But the same old question arises about what constitutes a random versus a non-random assignment as arises about appropriateness and relevance.

The solution I propose is based on the notion of a "plausible chain of reasoning." It is this. For an indirect speech act to be understood as being relevant, or appropriate, it must be able to participate as a premise, or as a sub-conclusion, in a quasi-logical, or better, psycho-logical, chain of reasoning that plausibly relates the event that initiates it to its intended illocutionary force. To see the full implications of this proposal, let us see how it works with the example. The man is offered some cake, and this offer is the event that initiates his response. From the perspective of the man, (9) is true.

- (9) I am being offered some cake.

Such an event calls for one of two responses, an acceptance or a refusal, appropriately modified by politeness conventions. Let us assume that "Yes, please" and "No, thank you" count as direct, literal speech acts for accepting and refusing, respectively. They certainly are conventionally regarded as direct ways of accepting and refusing. Now we can see that in this particular case, the proposal is this for (7) to be understood as a refusal, it must be able to participate as a premise, or as a subconclusion, in a psycho-logical chain of reasoning that plausibly relates

the original offer to its acceptance or rejection. Such a chain of reasoning might look something like (9) - (15).

- (9) I am being offered some cake
- (10) I am on a diet
- (11) People on diets ought not to eat fattening things
- (12) Cake is fattening
- (13) (It follows logically that) I ought not to eat any cake
- (14) (It follows deontically that) I will not eat any cake
- (15) (It follows conventionally that) I will refuse the cake

This chain of reasoning, including the intermediate and final conclusions does not constitute a deductively valid argument in the usual logical sense. The relationships that exist between (13) and (14), and between (14) and (15) are not entailment relations, but they are characteristic of human reasoning.

A number of important observations have to be made about the chain of reasoning--observations that amount to constraints on what it normally is. First, there are no unnecessary premises in it. Every premise is needed for the establishment of the first subconclusion, (13), which in its turn is needed for establishing the final conclusion. Second, although the order of the premises that are introduced from the speaker's general knowledge can be manipulated, the most natural order is one in which each premise invokes a concept that has been foregrounded (in the sense of Chafe, 1972) by the preceding one. If this were not the case, the possibility of introducing irrelevant premises would arise--a possibility that could serve no useful purpose in the present context. In fact, this constraint probably needs to be a little more liberal than I have described, but for the reasons I have indicated, something close to it needs to operate. Third, the conclusion of the chain contains the information appropriate for a direct response to the initiating event, an event that need not itself be a linguistic one (as it is in the present example). The initiating event might be an observed event to which an appropriate response might be a description of it, or of a reaction to it. Consequently, in the general case, "response" should not be taken to mean "reply."

There are doubtless other constraints that a more detailed analysis would reveal, but for the moment I want only to suggest that the conjunction of these (or some comparable set of) constraints constitutes what I mean by "plausibility" in the context of my requirement that the chain of reasoning be a psycho-logically plausible one.

We are now in a position to consider what happens from the perspective of the hearer. The most important thing is that the hearer assumes that the speaker is constrained in what he says in just the kind of way that I have indicated. The hearer, therefore, attributes to the speaker some plausible chain of reasoning. However, the hearer may not have all the knowledge that is available to the speaker (he may not know that he

is on a diet, for example). Consequently, he may have to make inferences of his own in order to reach some of the premises required. This would be true if, for example, the response to the offer of cake had been (12) rather than (10). Sometimes these inferences are incorrect and one understands correctly what was intended, but for the wrong reasons, or one misunderstands it altogether. As we shall see, this fact, that the hearer's interpretation is only probabilistically determined, has some important consequences for the speaker's selection of his utterance.

The most crucial claim that I wish to make about the chain of reasoning is this. Assuming that the speaker does not choose to express himself directly (for whatever reason), then within the limits of the context, any of the premises or subconclusions in the chain from the initiating event to the (direct) conclusion can function as more or less easily interpretable surrogates for the conclusion--any of the steps can constitute an indirect speech act appropriate to the direct speech act that constitutes the conclusion. Thus, any of (10), (11), (12), (13), and (14) can serve as indirect response to the offer. And, if some other response is made, it must be able to serve as a step in a similar chain of plausible reasoning. If it cannot, it is an inappropriate response. It is precisely these constraints that prevent (8) from being a possible indirect response to the offer, since there is no basis of shared knowledge that will normally permit a hearer to reconstruct an argument in which (8) figures to be relevant on the chain from initiating event to conclusion.

An important question that now needs to be answered is why do people use language indirectly in the first place, and why, given that they can choose from a restricted range of indirect communicative acts, do they select the ones they do. Why, for example, would a speaker choose (10) instead of, say (12)? The answer to the first part of the question depends on exactly what kind of indirect language act is being used. For example, metaphors may be used for purposes of communicative economy, communicative vividness, or even communicative possibility (see Ortony, 1975). With indirect speech acts, the answer is very often that the speaker gets "two for the price of one." For example, he can, with one utterance, not only refuse the offer, but also satisfy certain social conventions by providing a good reason for his refusal, or at least hinting at one. As Searle (1975) points out, in an indirect speech act the speaker intends both the sentence meaning and the speaker meaning to be recognized by the hearer. So, indirectness affords economy as well as, often, politeness and sensitivity.

There remains the question of why a speaker should select one form over another. The answer again lies in the fact that the communication of the literal meaning of the indirect language act is intended. Some of the knowledge that is needed to construct the reasoning chain may be more publicly available than other knowledge required. Thus, most people know that people on diets ought not to eat fattening things

(depending on the purpose of the diet, of course). Consequently it can be assumed that a hearer has more ready access to that fact than to the fact that the speaker is on a diet (which possibly very few people know). Thus, the speaker's selection of the particular language act can take advantage of his beliefs about what the hearer is likely to know. It can also take advantage of the fact that some of the choices seem to have a stronger force than others. This is a complex issue. My intuitions are that (13) leaves open the possibility of ultimately accepting some cake rather more readily than does (10), perhaps because once (10) is used it must be relevant to the chain of reasoning, whereas if (13) is used, it could be used to reach a different conclusion. After all, most people occasionally do things that they ought not to do, and that possibility seems wide open if the response to the offer is (13). This is not the place to explore these issues further, but it is worth noting that many jokes capitalize on expectations of plausible reasoning chains of the kind I have been discussing--the trick is to make them go awry!

My proposal shares certain characteristics with that of Searle (1975) in that it suggests a not necessarily conscious chain of reasoning. It differs from Searle's account insofar as it makes claims about the constraints on what can be said and understood. Searle's chain of reasoning contains many metalinguistic premises about indirectness that I have treated as background assumptions. My focus, by contrast, is on the content of the chain. What I have proposed is a possible answer to the question "How indirect can an indirect speech act be?" I have suggested that it cannot be so indirect that it could not participate in a chain of plausible reasoning relating a representation of the initiating event to an appropriate direct response to that event. I have also suggested that the illocutionary effect of all steps within such a chain will be appropriate for that initiating event.

Definite Descriptions of Inference

When speakers and writers produce, and hearers and readers comprehend definite descriptions, they do so against a background of knowledge that includes their tacit knowledge about indirectness. This knowledge is often brought to bear in dealing with definite descriptions of inference--descriptions, that is, in which the underlying predicate could appear on the reasoning chain and that could constitute a direct or indirect speech act. Thus, for example, (3) and (4) are cases in which the underlying predicate could constitute the conclusion of a chain of reasoning--i.e. a direct language act, while (5) is a case in which the underlying predicate could constitute a premise in a plausible chain of reasoning--i.e. an indirect language act.

In a sense, what I have proposed is a partial account of relevance in Grice's (1975) sense, or perhaps better yet, a partial account of when apparent violations of relevance are indeed only apparent, and why. It is quite clear that the predicates underlying definite descriptions have to be relevant to the discourse just as any other

comparable meaningful components of it must be. This is the sense in which I claim that the constraints that govern what definite descriptions can be used by a speaker who hopes to be understood are the same as those that govern what a speaker can in general say, if he has those same aspirations. There is no doubt that a detailed translation of my proposals about indirectness into comparable ones about definite descriptions is no easy matter. One reason is that the reasoning process that underlies the determination of a referent may be from a conclusion to an initiating event, as in (3). Another is that not all the steps in the chain can be employed, but only those that contain information applicable to the referent--not, for example, generalizations like (11). In such cases, if the premise is to be incorporated, it has to be embedded as the complement of an appropriate verb of propositional attitude. Nevertheless, it seems to me that some of the notions that I have laid out might prove helpful, if only by virtue of the fact that they may eventually lead to better proposals by others.

Earlier, I suggested that perhaps descriptions of inference involving indirect speech acts and those not involving them, really hinge on fundamentally the same kind of processes. On the surface, the basic difference concerns whether or not they involve the addition of pragmatic knowledge. It turns out, however, not to be an easy matter to decide what is pragmatic knowledge and what is merely semantic or factual. For example, (16) is a description of inference:

- (16) The navigator had heard that the weather might be unpleasant. He had always been concerned for the comfort of the passengers. He proposed taking a more indirect route to avoid the possible storms. The captain disagreed. He felt that the sooner they arrived at their destination the better--he wanted a drink and a decent meal. After a long argument the cautious one got his way.

Now, for a hearer to determine that the cautious one and the navigator are coreferential, it is necessary for him to invoke general knowledge about what constitutes a cautious act (contrast this with an offensive act). This in turn requires inferences to be made about human actions and intentions. To be sure, the actions in question are not linguistically performed acts, but that appears to be the only difference. Furthermore, had the pilot asserted that he wanted to take the shortest route because of his frivolous desires, would he not, thereby have been indirectly recommending a (possibly) reckless act? Surely, what is pragmatic and what is not cannot come down to performative verbs. Yet, if it is to be broader than that, what criteria are to be used to separate the semantic from the pragmatic? The old notion of semantics as entailment is certainly too restrictive to be useful as a model of natural language processing, but the new notion of pragmatics seems to amount to little more than the notion that language processors are rational beings who engage their reasoning processes in language comprehension and production

just as they do in perception and action. Even the notion of a speech act seems to have very fuzzy boundaries unless it is trivialized by invoking psychologically 'interesting surface structural aspects like the presence or absence of performative verbs.

Distinctions between different classes of linguistic phenomena are usually difficult to maintain in any rigid way, particularly if they are supposed to have psychological correlates. This is true of the distinction between syntax and semantics, of that between semantics and pragmatics, of that between literal and nonliteral, and of that between descriptions of entailment, and descriptions of inference. As usual, clear cases are easy to recognize, but there is always a large grey, undecided area in the middle where the classification seems sterile. In the case of the distinction between descriptions of entailment and descriptions of inference, the problem is exactly the same as the classical philosophical one that plagues the analytic/synthetic distinction. This is hardly an accident since my distinction is really no more than the analytic/synthetic distinction in disguise. Maybe all that needs to be said is that some inferences (e.g. ones based strictly on the rules of logic) are generally easier to make than others. If this is right then it merely means that some relationships between descriptions and their intended referents are more transparent than others. Nobody could object to that.

The last question I want to deal with is the psychological status of my claims, particularly with respect to the inference patterns that I have proposed. My position is not that it is a necessary condition for the comprehension or production of a definite description of inference that a person actually construct such a chain of reasoning. My claim is only that it should be possible to do so--there has to exist some determinable connection between the predicate underlying the definite description and the discourse in which the description occurs. But, being determinable and being determined are different things. As a matter of fact, there are often other clues that will permit the hearer to make a good guess about the referent's identity, discourse topic being one of them. It is almost certainly the case that people sometimes do go through some such reasoning process as I have outlined, and if and when they do not, they could probably be induced to do so by being asked suitable questions about what they took the referent to be, and why it was reasonable or plausible to do so.

References

- Austin, J. L. How to do things with words. Oxford: Oxford University Press, 1962.
- Chafe, W. L. Discourse structure and human knowledge. In J. B. Carroll and A. O. Freedle (Eds.), Language comprehension and the acquisition of knowledge. Washington, D.C.: V. H. Winston, 1972.
- Grice, H. P. Logic and conversation. In P. Cole and J. L. Morgan (Eds.), Syntax and semantics

(Vol. 3): Speech acts. New York: Academic Press, 1975.

Morgan, J. L. Two types of convention in indirect speech acts. In P. Cole (Ed.), Syntax and semantics (Vol. 9): Pragmatics. New York: Academic Press, 1978.

Ortony, A. Why metaphors are necessary and not just nice. Educational Theory, 1975, 25, 45-53.

Ortony, A., & Anderson, R. C. Definite descriptions and semantic memory. Cognitive Science, 1977, 1, 74-83.

Searle, J. R. Indirect speech acts. In P. Cole and J. L. Morgan (Eds.), Syntax and semantics (Vol. 3): Speech acts. New York: Academic Press, 1975.

Searle, J. R. Indirect speech acts. In A. Ortony (Ed.), Metaphor and thought. Cambridge: Cambridge University Press. in press.

Stalnaker, R. C. Pragmatics. In D. Davidson and G. Harman (Eds.), Semantics of natural language. Boston: D. Reidel Publishing Co., 1972.

Footnotes

¹This research was supported in part by the national Institute of Education under Contract No. US-NIE-C-400-76-0116, and by a Spencer Fellowship awarded to the author by the National Academy of Education. I am grateful to my colleagues Glenn Kleiman, Robert Kantor, and Jerry Morgan for their helpful comments on earlier drafts of this paper.

²I call them definite descriptions of entailment because technically they both are. The principle of identity, that $p \supset p$, represents an admittedly trivial entailment. It is important in the present context because it represents the case in which some predicate is literally transformed into the body of a definite description. More complex cases are still based on the usual rules of propositional logic such as modus ponens, $((p \supset q) \cdot q) \supset q$.

³Caution is needed here. Some cases of giving do not entail having. One can give somebody a pat on the back, or a kick in the teeth; the recipient gets it alright, but he doesn't have it! However, if we specify the appropriate constraints on the object the entailment will hold.