Reference Diaries¹

Herbert H. Clark
and
Catherine Marshall

Stanford University

When we make a definite reference to a thing, we normally make sure that our audience "shares" with us certain knowledge about that thing. To refer to a woman as she, the woman, or Nancy, we usually have good evidence that our audience knows about her too. But exactly what "shared" knowledge is required? This question is critical if we are ever to discover how people make or interpret definite reference--how they represent knowledge in memory and consult it in uttering and interpreting expressions like she, the woman, or Nancy. The question is critical if we are ever to characterize the mental archive people have for storing the facts they need to know for definite reference. We will argue that this archive has to be another detailed diary, or reference diary, supplemented by atlases, histories, and certain other reference texts. To make this argument, and to see what goes into the archive, we will examine the prior question, What "shared" knowledge is required for definite reference? As it happens, this question leads directly to a puzzle we will call the mutual knowledge paradox. It is in the solution of this puzzle that we get our best clues as to what the reference diary must be like.

The Mutual Knowledge Paradox

Imagine that there is a Marx brothers retrospective on at the local theater for which there are two or three movies a night for several evenings. Against this background consider the following scenario:

Version 1. On Wednesday night Ann and Bob go to see Monkey Business. The next morning Ann meets Bob and asks, "What did you think of the movie?"

What we are interested in is Ann's definite reference the movie, which she intends to refer to Monkey Business. What facts does Ann have to assure herself of before she can felicitously make this reference? Our interest here is in only those facts that are involved in "shared" knowledge. As a first condition, for example, Ann must herself have a certain awareness of Monkey Business. For now we will express that awareness as "knowing about R" (where R stands for the referent Monkey Business). Thus one fact Ann must assure herself of is this:

- (1) Ann knows about R.

But is this enough? Of course not, for (1) provides no assurance that Bob knows about Monkey Business. The way it fails can be made clear in a variation on the original scenario that goes like this:

Version 2: On Wednesday night Ann and Bob go to see Monkey Business, but neither knows that the other went too. The next morning Ann meets Bob and asks, "What did you think of the movie?"

Although version 2 satisfies condition (1), Ann has clearly made her definite reference without the right assurances. If (1) were all that had to be satisfied, version 2 would lead to a felicitous definite reference. Since it does not, we must add another condition, and the obvious one is that Ann must also assure herself that Bob knows about the movie, condition (2):

- (2) Ann knows that Bob knows about R.

(If it seems too strong to require knowledge instead of belief, each know can be replaced by believe; without legislating on the argument, we will stick with know).

At first, conditions (1) and (2) together seem enough, but it is easy to show that they are not. Consider this version of the original scenario:

Version 3: On Wednesday night Ann goes to see Monkey Business, and there she sees Bob. But he doesn't see her, and she realizes this. Furthermore, she realizes that Bob, unlike herself, might have seen A Day at the Races and A Night at the Opera, which are also showing that night. The next morning Ann meets Bob and asks, "What did you think of the movie?"

Although Ann has satisfied conditions (1) and (2)--she knows about Monkey Business and she knows that Bob knows about Monkey Business--she has not yet assured herself of enough. She cannot be sure Bob won't take the movie as referring to A Day at the Races or A Night at the Opera or even some other movie. Why? Because he couldn't be sure, uniquely, which movie she had in mind that he knew about. Bob must know not only about Monkey Business, but also that Ann knows about Monkey Business. At least, this is something Ann must try to assure herself of. This leads directly to the next

condition:

- (3) Ann knows that Bob knows that Ann knows about R.

With condition (3) we must surely have strong enough conditions for the success of Ann's definite reference. But that isn't so, as we can show in still another variation on the original scenario:

Version 4: On Wednesday night Ann goes to see Monkey Business, and there she sees Bob. As she walks down the aisle, she notices that he sees her, but as she is about to wave he turns and moves to another part of the theater. So she does not believe that he realizes that she has seen him. The next morning Ann meets Bob and asks, "What did you think of the movie?"

This version satisfies conditions (1), (2), and (3). Ann knows about Monkey Business; she knows that Bob knows about it; and she knows that he knows that she knows about it. But Ann doesn't believe that he knows that she knows that he knows about it. This piece of negative knowledge should be enough to keep Ann from using her definite reference. What if Bob had gone to A Day at the Races and A Night at the Opera too?, she should ask herself. He might think that while he is sure she didn't see him at Monkey Business, she might have seen him at one of the other two. If so, she might be referring to one of the other two. He couldn't be sure. According to Ann's reasoning, therefore, she must assure herself of something more--that Bob realizes that Ann realizes that he had been to see Monkey Business. That is, she must satisfy the following condition too:

- (4) Ann knows that Bob knows that Ann knows that Bob knows about R.

With condition (4) it looks as if we have gone far enough (see Kempson, 1975, p. 165; Stalnaker, 1977, p. 137), but can we be sure? Only if we can't dream up another variation that satisfies conditions (1) through (4) but still doesn't work. Indeed, with a little difficulty, we can:

Version 5: On Wednesday night Ann goes to see Monkey Business and there she sees Bob and Charles. Because she sits down a few rows in front of them, she believes that they see her there, but because she doesn't turn around while they are there, she believes that they don't realize that she has realized that they have seen her there. On the way home, however, she meets Charles, who tells her that Bob did realize that she had seen them there, but because she hadn't waved at them, Bob was certain that she didn't realize that they had seen her notice that they were there too. The next morning Ann meets Bob and asks him, "What did you think of the movie?"

Complicated as this version is, we realize that Ann in good conscience shouldn't have made this definite reference. Although conditions (1) through (4) are all satisfied, Ann should have considered this possible reasoning on Bob's part. What if Bob had seen A Day at the Races and A

Night at the Opera too. He might think that she had seen him at, say, A Day at the Races and that she thought he had seen her there too, even though he hadn't. He would then have reason to think she was referring to A Day at the Races, since to have referred to Monkey Business she would have been sure that he knew that she knew that he knew that she was there (Bob's equivalent to condition (4)). So despite all of the conditions she has already assured herself of, she must add one more:

- (5) Ann knows that Bob knows that Ann knows that Bob knows that Ann knows about R.

Is condition (5) enough? Hardly. What these versions show is that there is a way in principle of demonstrating that the last piece of iterated knowledge is insufficient. The method is this. Corresponding to Ann's condition (1) is an analogous condition that Bob must assure himself of if he is to uniquely identify the referent for Ann's definite reference, and it is this:

- (1') Bob knows about R.

For Ann to be sure that her reference goes through, she must put herself in Bob's shoes, reason as Bob would, and make sure that he would identify the intended referent uniquely. What we did in constructing version 2 was create a scenario in which (1) and (1') held, but Ann couldn't know that (1') held. This led us to add condition (2), Ann knows that Bob knows about R, the equivalent of Ann knows that (1'). But just as Ann needs to assure herself of (2), Bob needs to assure himself of (2'):

- (2') Bob knows that Ann knows about R.

But then (2') is something else Ann needs to know, as we showed in creating version 3 of our scenario, and this led to condition (3). Corresponding to (3), however, is Bob's (3'), which we used in creating version 4. In principle, we could use this procedure to construct countermanding versions ad infinitum.

The successive versions and the conditions they give rise to eventually become absurdly complicated, but they do bring out a general point. In principle, one must satisfy oneself of an infinite number of conditions either to make or to interpret a definite reference. Hence the mutual knowledge paradox. If each condition takes a finite amount of time to check, no matter how small, and if these checks cannot all be made in parallel, then making or interpreting a definite reference like the movie should take an infinite amount of time.

Mutual Knowledge

In common parlance, "shared knowledge" has several definitions. Ask your aunt what it means for the two of you to share knowledge that the mayor is an embezzler, and she would probably say, "It means that you know he is an embezzler, and that I do too." If we express the proposition that the mayor is an embezzler as p, then the first definition of shared knowledge comes out like this:

A and B share₁ knowledge that p = def.

(1) A knows that p.

(1') B knows that p.

However, she might give you a more complicated answer: "It means that both of us know that he is an embezzler, and furthermore, I know that you know he is, and you know that I know he is." This leads us to a second definition of shared knowledge:

A and B share₂ knowledge that p = def.

(1) A knows that p.

(1') B knows that p.

(2) A knows that B knows that p.

(2') B knows that A knows that p.

Indeed, we can define a series of types of "shared" knowledge merely by extending the list of statements. None of these finite definitions, of course, describes the "shared" knowledge required of Ann and Bob in her reference to Monkey Business. For that we need something more.

What is required, apparently, is the technical notion of mutual knowledge. It has been defined and exploited by Lewis (1969) and Schiffer (1972) for dealing with close cousins of the problem we have raised here. Mutual knowledge is Schiffer's term, while Lewis' term for the same thing is common knowledge. We have opted for Schiffer's term since it seems more transparent and less open to misinterpretation. In any case, mutual knowledge is defined as follows:

A and B mutually know that p = def.

(1) A knows that p.

(1') B knows that p.

(2) A knows that B knows that p.

(2') B knows that A knows that p.

(3) A knows that B knows that A knows that p.

(3') B knows that A knows that B knows that p.

et cetera ad infinitum.

Heuristics for Assessing Mutual Knowledge

So far two conclusions seem firm. First, definite reference requires a certain amount of mutual knowledge. Other simpler notions of "shared" knowledge will not do. Second, it is unthinkable that speakers and listeners assess mutual knowledge by working serially, statement by statement, through the infinity of statements that make up mutual knowledge. But they surely assess it somehow, as the first conclusion seems to require. The inevitable conclusion is that they use some sort of heuristics. We will consider two families of such heuristics--truncation heuristics and co-presence

heuristics.

Truncation Heuristics

The stickler in assessing mutual knowledge statements is that there is an infinity of such statements, and that is too many to check. What if people checked only a few of them--like the first four? The task could then be carried out in a finite, even short, amount of time. There would be errors, of course, but they would probably be neither very serious nor very frequent. If Ann has verified the statement (4), Ann knows that Bob knows that Ann knows that Bob knows that p, it is extremely likely, on actuarial grounds, that the higher order statements would check out too. And when she does make an error, Bob will often look puzzled or ask for clarification, which will allow her to repair her reference. Indeed, repairs are quite frequent in spontaneous speech as if speakers might be doing just that. So people could assess only a truncated part of mutual knowledge. Heuristics of this kind will be called truncation heuristics.

Are these heuristics plausible as the way people normally assess mutual knowledge? We believe not. Our doubts lie in two areas. First, it is not easy to deal with statements as complicated as (4). It is implausible that people check these statements per se. Second, the evidence needed to verify such statements anyway suggests a radically different family of heuristics.

In version 4 of our movie scenario, Ann didn't believe that Bob knew that she knew that he knew about Monkey Business, a violation of knowledge statement (4). Version 4 is complicated. Not only did we have a hard time creating it, but people have a hard time grasping it, for it is difficult to keep track of who knows what. Statements like (4) are difficult not because of their syntactic form, but because they describe reciprocal relations between two people. Whereas John Dean knew that Nixon knew that Haldeman knew that Magruder knew that McCord had burgled O'Brien's office is fairly comprehensible, John Dean knew that Nixon knew that John Dean knew that Nixon knew that McCord had burgled O'Brien's office is not. Although when we need to we can figure out fourth order reciprocal relations--not just the statements themselves, it seems highly implausible that we do so routinely.

But what counts as evidence for the truth of statements like (1), (2), (3), and (4)? Take statement (3), Ann knows that Bob knows that Ann knows about R. Obviously, Ann won't have this statement per se already stored in memory. She doesn't go through life creating statements like this for every object she or anyone else might want to refer to. Rather, what she needs to verify (3) is a piece of evidence from which she can deductively or inductively infer it. Imagine that she and Bob had gone to Monkey Business together. It is hard to think of better evidence than this that she could appeal to for the truth of (3). Of course, the inductive rules by which she infers (3) from this evidence need to be spelled out, but that doesn't sound impossible.

The fact that Ann and Bob saw the movie together, however, is more useful evidence even than that. It is also about the best evidence we could imagine for the truth of (1), and of (2), and of (4), and so on ad infinitum. It is a piece of evidence that allows Ann, in one quick jump, to be sure of the truth of all the statements. Why, then, would she want to check the statements one by one--even a truncated list of them? She would be better off looking for that single piece of evidence that could in principle confirm them all. Indeed, that is the foundation assumption of the next family of heuristics we will take up, the co-presence heuristics.

Consider the following strategy. When people make or interpret a definite reference, they try to assure themselves of mutual knowledge of the referent by searching for evidence of what we will call triple co-presence. This is evidence of a particular event in which the speaker, listener, and referent are "co-present," i.e., are "present" simultaneously, as when Ann, Bob, and Monkey Business are openly "present" together on Wednesday night. Strategies like this will be called co-presence heuristics. To see how they are reasonable, we will look at first principles.

When Lewis and Schiffer hit on the notion of mutual knowledge, both recognized the need for a finite means of handling the infinity of statements. Their solutions were essentially the same. If A and B make certain assumptions about each other's rationality, they can use certain kinds of evidence, or states of affairs, to infer that each one of the infinite number of statements in mutual knowledge is true. But how? We get some hints from a concrete illustration of mutual knowledge devised by Schiffer.

The scene: Ann and Bob are sitting across a table from each other, and there is a single candle between them. Both are looking at the candle, and both see the other looking at it too. The proposition *p* is that there is a candle on the table. Consider the scene from Ann's point of view. Clearly, she has direct evidence for the truth of (1):

(1) Ann knows that *p*.

But she also sees that Bob has his eyes open and is looking simultaneously at her and the candle. That is, she has evidence that she and Bob are looking at each other and the candle simultaneously. We will call this the simultaneity assumption. Indeed, she assumes that he is not only looking at her and the candle, but also attending to them. We will call this the attention assumption. Finally, she assumes that Bob is normal and in her shoes he would be drawing the same conclusions she is. We will call this the rationality assumption. But if Bob is attending to the candle and is rational, he has evidence for (1'):

(1') Bob knows that *p*.

This, however, is Ann's conclusion, and so she has evidence for (2):

(2) Ann knows that Bob knows that *p*.

But if Bob is rational, he will be drawing the inference that corresponds to hers--his equivalent of (2)--namely (2'):

(2') Bob knows that Ann knows that *p*.

Once again, this is Ann's conclusion, and so she has evidence for (3):

(3) Ann knows that Bob knows that Ann knows that *p*.

In like fashion, Ann would be justified in iterating this process through the remaining knowledge statements (4) through infinity, and Bob would be justified in doing the same for his.

So Ann has reason to believe that she and Bob mutually know that there is a candle on the table. First, there is the "direct" evidence. She directly perceives that there is a candle on the table and that Bob is simultaneously looking at both her and the candle. Second, there are her assumptions about the situation. She assumes that Bob is consciously attending to her and the candle, that he is doing so at the same time she is, and that he is rational. The upshot is that she has no reason to believe that she couldn't confirm the knowledge statements as far down the list as she wanted to go. She is therefore justified in claiming mutual knowledge. Indeed, since nothing she doesn't know herself can be mutual knowledge, and since she can assume Bob is chronically rational, all she needs to do normally is search for evidence of her and Bob simultaneously attending to each other and the candle on the table. With this we have the essence of the co-presence heuristics: To assess mutual knowledge, people search for evidence of triple co-presence--an event in which A and B are simultaneously attending to each other noting the same evidence for *p*. In equation form:

Co-presence + Assumptions = Mutual knowledge

The co-presence heuristics both solve the mutual knowledge paradox and make intuitive sense. When we assure ourselves of mutual knowledge, it is unlikely that we check for a series of pieces of evidence, even as few as the truncation heuristics might let us get away with. More likely, we check for a single piece of evidence of just the right kind. The candle example suggests that what we check for is evidence of triple co-presence.

Varieties of Triple Co-presence

There are many different kinds of evidence people may use for the triple co-presence of the speaker, listener, and referent. Some of these constitute strong evidence for triple co-presence, and others constitute weak evidence. That is, some kinds rightly give people a lot of confidence that the referent is mutually known, whereas other kinds do not. As reflected in our equation, there is a trade-off between the evidence and the assumptions. The stronger the evidence is, the fewer assumptions are needed to infer mutual knowledge. Conversely, the fewer assumptions that are needed, the stronger the evidence is considered to be. The strongest evidence requires the fewest, or weakest, assumptions.

The cornerstone of our argument is this. The prototypical kind of evidence for mutual knowledge is physical co-presence, very much as illustrated in Schiffer's candle example. It is the strongest possible evidence, the one requiring the fewest auxiliary assumptions, and all other kinds are weaker in one way or another. What follows is a tentative classification of these varieties of triple co-presence.

1. Physical co-presence. Ann, Bob, and the candle are an example par excellence of physical co-presence. Not only are the three of them physically present together, but Ann can readily assume that Bob is attending to this fact, is doing so at the same time she is, and is rational. All three auxiliary assumptions are necessary. If she believed Bob was catatonic, or hypnotized the right way, or very near-sighted, for example, she wouldn't want to assume physical co-presence. Once Ann has assured herself of the direct evidence and these assumptions, she is warranted in inferring mutual knowledge of the candle and can refer to it as the candle.

There are two distinct types of physical co-presence. Ann may refer to the candle while it is still physically co-present with them, as in The candle is romantic, isn't it? Or she may refer to the candle some time after it has been co-present with them, as in The candle was romantic, wasn't it? These two types could be called immediate and delayed physical co-presence. The first kind, on the face of it, is the stronger evidence. When physical co-presence is synchronous with the definite reference, Ann can be sure that she and Bob mutually know about the candle at the time she is referring to it. She doesn't have to count on Bob's remembering the past incident of physical co-presence, as she does in the delayed kind.

The assumptions Ann would need in order to infer mutual knowledge from immediate physical co-presence are these: simultaneity, attention, and rationality. She would need an additional one for the delayed case: simultaneity, attention, rationality, and memory. Simultaneity, attention, and rationality refer to the assumptions we have described earlier. Memory refers to the additional assumption for delayed physical co-presence: Ann must assume that Bob can and will recall the earlier incident of their physical co-presence. So far so good. The stronger the evidence, the fewer assumptions Ann needs in order to make her definite reference. Immediate physical co-presence has one fewer requirement than delayed physical co-presence.

2. Linguistic co-presence. Many things we refer to have never been physically co-present. They are often things we or someone else has mentioned in conversation. Imagine Ann saying to Bob I bought a candle yesterday. Her utterance of a candle is a locutionary act that posits for Bob the existence of a particular candle in the real world. If Bob hears and understands a candle correctly, he knows about the candle's existence at the very same time as she posits it. It is as if Ann has placed the candle on the stage before the two of them so that it would be physically co-present. So when Ann utters a candle and Bob simultaneously

understands it, the two of them can be said to be in the linguistic co-presence of the candle. Once Ann has established this, of course, she can make a definite reference to it, as in The candle cost me plenty.

Linguistic co-presence is weaker evidence for mutual knowledge than physical co-presence. Seeing is believing--hearing about something isn't. To begin with, linguistic co-presence requires the assumptions of simultaneity, attention, and rationality. Ann and Bob must be attending to Ann's utterance of a candle simultaneously, and both must be rational. And like delayed physical co-presence, linguistic co-presence requires memory. For Ann to refer to the candle, she has to count on Bob's recalling the earlier incident of linguistic co-presence with her uttering of a candle. But there is an additional assumption we will call understandability. Ann must assume that Bob will penetrate her indefinite reference, a candle, and understand that she is sincerely positing the candle's existence. She must assume that Bob understands her, and he must assume that she believes he does.

3. Indirect co-presence. Imagine Ann saying to Bob I bought a candle yesterday; the wick is made of cotton. In uttering a candle Ann has established the linguistic co-presence of him, her, and the candle, but not of him, her, and the wick. How, then, can she refer to the wick? She has to assume that when Bob accepts the existence of the candle, he will also accept the existence of its wick. This way, by uttering a candle, Ann has established what we will call the indirect co-presence of her, Bob, and the wick.

The inferences required in indirect co-presence are often much stronger than those needed for wick (see Clark, 1977; Clark & Haviland, 1977). Ann can refer to something that is only likely to be associated with a thing she has already established, or even only possibly associated with it. She can tell Bob: I bought a candle yesterday, but the wrapper was torn; or I bought a candle yesterday, and the bayberry smelled great. Candles don't necessarily come in wrappers nor are they often made of bayberries, yet these are parts she expects Bob to infer on the basis of her definite references to them. So what is established may be only the likelihood or possibility of a thing being co-present with the speaker and listener. Its certain existence is established only with the definite reference itself.

Indirect co-presence is parasitic. It has to be established via some other type of co-presence--for example, physical or linguistic co-presence. Before Ann can say The price was \$3 of a candle, she must already have established the candle's co-presence. She and Bob could be looking at it, for physical co-presence, or she could have just mentioned it, for linguistic co-presence. For the moment we will assume that indirect co-presence is always established via either physical or linguistic co-presence.

There is both a strong and a weak case of indirect co-presence. Instead of saying The price was \$3, Ann could have said The price of the candle

was \$3, providing a much more certain reference. She would have made it explicit that the price referred to is that of the candle and not of something else. Bob would then have had no trouble inferring that there was one and only one price associated with the candle. They both could then assume that they mutually knew about the price. This case may be so direct that it ought to be placed in a separate category. For now we will treat it as a very strong kind of indirect co-presence.

To infer mutual knowledge from indirect co-presence, Ann and Bob need all the assumptions of physical or linguistic co-presence, whichever is the parasite's host, plus one we will call associativity. They have to assume that each other is capable of entertaining the certainty, likelihood, or possibility of a particular part or role being associated with the thing whose co-presence has already been established. The hierarchy still works as expected. Indirect co-presence, because of its added assumption, is weaker evidence for mutual knowledge than either physical or linguistic co-presence.

4. Cultural co-presence. Even when Ann is not acquainted with Bob, she can assume there are particulars the two of them mutually know. The basic idea is that there are things everyone in a culture knows about. She reads newspapers, and so does everyone else in her culture. So Bob and she can mutually assume that they both read newspapers. Ann can then take the fact that John Dean, Michael Doonesbury, and Billy Jean King have been prominently mentioned in the newspaper as good evidence that she and Bob mutually know about these people. This is an instance of what we will call cultural co-presence. Certain particulars are assumed to be universally known in a cultural milieu--they are culturally co-present for everyone in it--and that is taken as evidence that everyone in the milieu knows about them.

The trick, of course, is to judge cultural milieus. Ann may think that she and Bob mutually realize that they are both high school graduates, or drug dealers, or nineteenth century history buffs, or New Yorkers, or telephone operators, or some combination of these, and her assumptions about cultural co-presence will change accordingly. If her assessments are accurate, her definite reference is likely to succeed, and if not, it isn't.

Cultural co-presence doesn't appear to belong to the same hierarchy as the previous three types of co-presence. For one thing, it is relatively permanent, whereas the other three are relatively transitory. Culturally known particulars take time to become familiar and to lose familiarity. Teddy Roosevelt is familiar to Americans today, just as he was 75 years ago. Particulars known by physical, linguistic, or indirect co-presence have only fleeting familiarity and then only to specific pairs of people. Mutual knowledge about these particulars is easily established, but also easily lost. For another thing, cultural co-presence is parasitic on other forms of mutual knowledge. For Ann to establish that she and Bob mutually know that they belong to the same cultural subgroup,

she must find evidence of triple co-presence of that fact. She might establish it, for example, through linguistic co-presence, as in, What do you know--we're both New Yorkers.

To infer mutual knowledge from cultural co-presence, therefore, people need assumptions that are not comparable with those of the other three types. Take Ann's reference to Hoover Tower in a conversation with Bob. First, she must assume that she and Bob mutually know that they belong to a particular cultural subgroup, say Stanford University students. We will call this assumption cultural membership. How Ann justifies this assumption, however, will not be simple. Like other types of mutual knowledge, it must be based on evidence of some kind of co-presence. Second, she must assume that virtually everyone in this cultural milieu takes it for granted that they all know about Hoover Tower. We will call this assumption universality of knowledge. The paucity of these assumptions should not fool us into thinking that cultural co-presence is strong, for they hide a tangle of complex justifications based on other pieces of evidence and other assumptions. It is best to treat cultural co-presence as incommensurate with the other three.

With cultural co-presence we have come to the last of the major kinds of co-presence. Not every kind of evidence for mutual knowledge, however, can be neatly classified as one of these four types. Some appear to require a complex combination of them, and not surprisingly, they provide intuitively weaker evidence for mutual knowledge.

Reference Diaries

If people assess mutual knowledge via triple co-presence, they must have a memory full of facts about triple co-presence. What do these facts look like? How are they represented? How are they assessed? If mutual knowledge is critical to definite reference--as we have suggested--then questions like these ought to be central to any theory of speaking, listening, or memory. Indeed, the arguments we have offered lead to a rather provocative conception of memory representation and memory search. It is provocative in that some of its critical properties are absent from most current models of comprehension and memory.

Most investigators have assumed that in processing definite reference people search memory for the particulars actually referred to. Take Ann's reference to Monkey Business. On hearing this Bob would search memory for a referential index to the intended referent Monkey Business. This index is a stand-in, so to speak, for the movie itself. Although the current models of comprehension differ in their specifics, virtually all of them assume this kind of search for the intended referent. That includes Anderson (1976), Clark and Haviland (1977), Kintsch (1974), Rumelhart, Lindsay, and Norman (1972), Schank and Abelson (1977), and Winograd (1972), to name just a few.

But if people use some kind of co-presence heuristics, then all of these models are incorrect--or at least incomplete. The point is that

Bob cannot search memory for the referent alone. That would hardly guarantee that it was mutually known to him and Ann, as it must be for her reference to be legitimate. Rather, he must search for an event that involves not only the referent but also Ann and him. That is, it must be an event of triple co-presence--of physical, linguistic, indirect, or cultural co-presence, or of some combination of these four types. In none of the current models just mentioned does the listener search for such an event.

Previous models of comprehension have treated search through memory as if it were a search through a telephone book. In a definite reference like the man in the red shirt we are told the name and address of the individual we want to get hold of. Our task is to search the telephone book for his number, our direct connection to him, his referential index. With the co-presence heuristics, memory must be more like a diary, more like the personal log Nixon kept of everything he did and experienced during his years at the White House. As before, in the man in the red shirt we are told the name and address of the individual we want to get hold of. But to find him we must search our diary for an entry that provides evidence of the co-presence of the speaker (say, Gertrude), us, and an individual of that description. The diary entry must show that we were physically or linguistically co-present, or that we were co-present in some other sense. That is, we must search in every case for an event. This is far more complicated than searching the telephone book, with or without yellow pages, for the right number.

The diary, of course, cannot be used alone. We also need histories and atlases to refer to John Dean, the Second World War, the decline and fall of the Roman Empire, and China, particulars that are culturally co-present. And for indirect co-presence we will also need texts on science, medicine, engineering, and law. To know that candles have wicks we need to look up facts about the engineering of candles.

What we need, in summary, is a diary of the significant events in our own personal experience, supplemented by cultural histories and atlases for cultural co-presence and by various reference texts for indirect co-presence. Such a diary contains a record of the events we will need for assessing co-presence. Anything less than a diary will be too little.

Footnote

¹This paper is an abbreviated version of "Definite reference and mutual knowledge," presented at the Sloan Workshop on Computational Aspects of Linguistic Structure and Discourse Setting, University of Pennsylvania, May 1978. We thank Eve V. Clark for her helpful comments on the manuscript.

References

- Anderson, J.R. Language, memory and thought. Hillsdale, N.J.: Erlbaum, 1976.
- Clark, H. H., & Haviland, S. E. Comprehension and the given-new contract. In R. O. Freedle (Ed.), Discourse production and comprehension. Norwood, N.J.: Ablex Publishing, 1977, pp. 1-40.
- Clark, H. H. Inferences in comprehension. In D. LaBerge & S. J. Samuels (Eds.), Basic processes in reading: Perception and comprehension. Hillsdale, N.J.: Erlbaum, 1977.
- Clark, H. H., & Marshall, C. Definite reference and mutual knowledge. Paper presented at the Sloan Workshop on Computational Aspects of Linguistic Structure and Discourse Setting, University of Pennsylvania, May 1978.
- Kempson, R. M. Presupposition and the delimitation of semantics. Cambridge: Cambridge University Press, 1975.
- Kintsch, W. The representation of meaning in memory. Hillsdale, N.J.: Erlbaum, 1974.
- Lewis, D. K. Convention: A philosophical study. Cambridge, Mass.: Harvard University Press, 1969.
- Rumelhart, D. E., Lindsay, P. H., & Norman, D. A. A process model for long-term memory. In E. Tulving & W. Donaldson (Eds.), Organization of memory. New York: Academic Press, 1972, pp. 197-246.
- Schank, R., & Abelson, R. Scripts, plans, goals and understanding: An inquiry into human knowledge structures. Hillsdale, N.J.: Erlbaum, 1977.
- Schiffer, S. R. Meaning. Oxford: Oxford University Press, 1972.
- Stalnaker, R. Pragmatic presuppositions. In A. Rogers, B. Wall, and J. P. Murphy (Eds.), Proceedings of the Texas conference on performatives, presuppositions, and implicatures. Arlington, VA: Center for Applied Linguistics, 1977, pp. 135-148.
- Winograd, T. Understanding natural language. New York: Academic Press, 1972.