

AUDIENCE DESIGN IN MEANING AND REFERENCE

Herbert H. Clark and Gregory L. Murphy

Department of Psychology
Stanford University
Stanford, California
U.S.A.

We argue that the speaker designs each utterance for specific listeners, and they, in turn, make essential use of this fact in understanding that utterance. We call this property of utterances audience design. Often listeners can come to a unique interpretation for an utterance only if they assume that the speaker designed it just so that they could come to that interpretation uniquely. We illustrate reasoning from audience design in the understanding of definite reference, anaphora, and word meaning, and we offer evidence that listeners actually reason this way. We conclude that audience design must play a central role in any adequate theory of understanding.

INTRODUCTION

In ordinary conversations, we tailor what we say to the particular people we are talking to. We have a good idea of the knowledge and beliefs they share with us at the moment and what they are thinking of, and we design our utterances accordingly. Evidence of this tailoring is everywhere. When we think our addressees share with us the knowledge that a man is named Aloysius, we may refer to him as Aloysius confident that they will understand who we are referring to. When someone in a conversation says I'm getting tired, we may reply with the highly elliptical So am I, confident that the others in the conversation will understand us. When we think our addressees share knowledge that a particular woman had just been sitting in a chair, we may point at the empty chair and say That woman is named Veronica and be confident that they will understand who we are referring to. We realize that anyone who didn't share our knowledge of Aloysius's name, or of the previous utterance, or of the person who was just sitting in the chair wouldn't necessarily understand us. We don't expect our utterances to be understandable by just anyone. They are intended for particular listeners with particular momentary thoughts and beliefs. Let us call this feature of utterances audience design (see Clark and Carlson, 1982a).

Although audience design is an obvious feature of language in use, it has rarely played a role in psychological models of language use. For years, the dominant model of understanding has focused on how we analyze sentences--that is, how we identify phonemes, words, syntactic constructions, and word meanings (see, for example, Carroll & Bever, 1976; Forster, 1979; Garrett, 1978; Marslen-Wilson & Tyler, 1980; Swinney, 1979). But when we listen to people in conversations, our aim isn't simply to identify the phonemes, words, and sentences they used. We try to understand what they meant in saying what they did on that occasion. For that, we must consider not only the acoustic signal and our knowledge of the language, but also the particular beliefs and thoughts we think the speaker shares with us at that moment. The dominant model, in effect, excludes the thoughts and beliefs of individual speakers and listeners. It is a model of understanding that excludes the understanders.

The dominant model of understanding would be worth pursuing if we could assume that people identify phonemes, words, and sentences without regard for audience design. But we cannot. Although some processes in understanding may not be greatly affected by audience design (see Marslen-Wilson & Tyler, 1980; Swinney, 1979), many others probably are. These include processes that have often been assumed to run off without regard for audience design. In this paper, we will review the role of audience design, especially as it has arisen in our own research, in three areas of comprehension--definite reference, anaphora, and word meaning.

DEFINITE REFERENCE

Definite reference has long been an object of study in philosophy and linguistics. In 1905, Bertrand Russell analyzed The king of France is wise as meaning: (a) there is a king of France; (b) there is not more than one king of France; and (c) there is nothing which is king of France and is not wise. Does this analysis do justice to people's uses of definite descriptions? Strawson (1950) argued that it didn't and proposed instead that when a person utters The king of France is wise, he presupposes (a) and (b) and asserts only (c). These two views and their offspring have been examined in detail by linguists and philosophers over the years (see Hawkins, 1978; Kempson, 1975; Wilson, 1975), and the debate still goes on.

For psychologists, the debate has an added dimension, since any model of understanding must say how definite reference is processed. In one influential model, Olson (1970) claimed that people view the object they are trying to refer to--the referent--as belonging to a "referent array," which may be either explicit and visible, or implicit. The speaker's reference to an object is successful if it describes the object so that it can be uniquely distinguished from all the other objects in the referent array. If a speaker wants to refer to one of five men at a bar, he cannot use simply the man at the bar. He must give a unique description for this array, like the tallest man at the bar.

Olson's model, however, is either incomplete or incorrect. First, for certain arrays, people give more specific descriptions than the model predicts. For an array with a dog and a rock in it, people are more likely to say Look at the dog than Look at the animal, or Look at the living thing, as Olson's model would predict (Cruse, 1977; Rosch, 1978). Second, references in conversations are dynamic. A person might first refer to someone as the woman who sold me a bottle of Chanel No. 5 this morning, but afterwards refer to her simply as the perfume lady, with only an allusion to the event described in the first reference (see Krauss & Weinheimer, 1966, 1967). And third, the model doesn't say how referent arrays get defined. This is an especially important issue when the array is implicit in the context, as it usually is in definite reference. When a wife tells her husband The door is open, what referent array is she assuming, and how does he know? For these and other gaps in Olson's account, there are no easy remedies.

From our point of view, the main problem is that Olson's model doesn't make full use of the audience design of a definite reference. As Clark and Marshall (1981) have argued, an essential part of audience design in definite reference is the use of the speaker's and addressee's mutual knowledge, beliefs, and suppositions, or what we will call common ground (see Clark & Carlson, 1981, 1982b). Imagine Anne talking to Pierre. To use the definite description the man you met yesterday, she has to have good reason to believe that Pierre can readily identify the referent uniquely on the basis of their common ground. Furthermore, Pierre should believe that Anne believes this. Pierre might have met many people the day before, but he knows that Anne doesn't know about most of these people. The only person they share knowledge of his meeting the day before is, say, Jacques, and so, his tacit reasoning goes, she must be referring to Jacques.

One premise in Pierre's tacit reasoning is especially important to that process. He knows there are many people who fit Anne's description of being a man that he had met the day before and, indeed, that Anne must assume there are too. But since he can assume that Anne has good reason for believing that he can readily identify the right person uniquely on the basis of their common ground, he can reason: "If she thinks I can readily identify her referent uniquely, she must be referring to the only person we mutually know that I met yesterday, namely Jacques." It is as if Anne's definite description weren't simply the man you met yesterday, but the man you met yesterday that I have good reason to believe that you can readily identify uniquely on the basis of our common ground. Implicitly, Anne's definite description includes an instruction to Pierre about how to reason through to the referent. Let us call this added premise the design assumption. As we will argue, in many cases a listener cannot identify the referent without making essential use of the design assumption.

How do speakers and addressees decode what is part of their common ground and what isn't? Since people can't look into each other's heads, they must use certain heuristics (Clark & Marshall, 1981). The heuristics they rely on work from three main sources of information and their combinations:

(a) Community membership. Once Anne and Pierre mutually recognize that they are both Parisians, they can take as common ground anything that is universally known, believed, or supposed by Parisians. Anne and Pierre, of course, may jointly belong to many other communities and subcommunities too, ranging, for example, from the community of all humans to the community of experimental psychologists who had studied with Piaget in 1959. Each of these communities and subcommunities has a body of knowledge, beliefs, and suppositions that everyone in the community takes to be universal in that community.

(b) Physical co-presence. Once Anne and Pierre have both experienced something together, like the seeing of a movie, or person, or scene, they can also each take what they experienced to be part of their common ground.

(c) Linguistic co-presence. Once Anne and Pierre have both participated in a conversation in which a fact has been asserted or an object mentioned, they can each take knowledge of the fact or the object to be part of their common ground.

So to design and understand references correctly, people need to consult the personal diary they keep in memory of the events in which they and others have taken part. For Anne to be able to use the man you met last night to refer to Jacques, she must assure herself, for example, that she and Pierre had been together when she had seen him meet Jacques. Without such a remembered event or some other basis for common ground, she can't be sure he will realize that Jacques is part of their common ground.

A striking example of reasoning from the design assumption can be found in demonstrative reference. Imagine Anne pointing at two men and saying to Pierre:

(1) Look at that man.

Under Olson's model, this reference should always be inappropriate, since listeners can never resolve the ambiguity: that man can always refer to either man. But imagine that one of the two men was distinctive in some way--say, he was dressed like Napoleon. Then Anne's reference would be quite appropriate, and Pierre would take her as referring to the man dressed like Napoleon. How? Reasoning from the design assumption, he would note that since she expected him to recognize which man she was referring to, there must be something distinctive about the referent to allow him to do that, and so she must be referring to the Napoleon look-alike.

The effectiveness of such reasoning was demonstrated in an experiment by Clark,

Schreuder, and Buttrick (in preparation). Buttrick walked up to students on the Stanford University campus and handed them a photograph of President Ronald Reagan standing next to his director of the budget David Stockman. People in an independent survey had said that they assumed that Reagan was well known to everyone but Stockman was not. Buttrick then asked each student one of two questions:

(2) You know who this man is, don't you?

(3) Do you have any idea at all who this man is?

Note that he used the same definite description, this man, in both utterances. Nevertheless, for 2, people offered replies such as "Yes, that's Reagan," not once thinking that he had referred to Stockman. For 3, people said such things as "Yes, I believe that's Stockman"; only two of 20 people assumed Buttrick was referring to Reagan. The reasoning people used relied on the presuppositions expressed in 2 and 3. In 2, Buttrick presupposed that they ought to know who he was referring to, but in 3; that he doubted they would. It was these presuppositions, and the community knowledge that Reagan is better known, that enabled them to come to a unique referent via the design assumption.

Audience design also plays a role in the level of description in definite reference. Recall that one of the facts Olson's model couldn't account for was that when people refer to objects in neutral contexts, they tend to describe them at a middle level of abstraction. When people refer to a dog, a bird, or a chair, they generally use the dog, the bird, and the chair in preference to more specific descriptions like the spaniel, the sparrow, and the armchair, and in preference to more general descriptions like the object, the animal, and the piece of furniture. This level of preferred naming has generally been called the basic level of abstraction (Rosch et al., 1976).

Why should the basic level be the preferred level of description? One explanation is that dog, bird, and chair denote basic categories, and these categories are basic because they maximize two factors simultaneously: (a) they are specific; and (b) they are dissimilar to other categories. The most general categories, like objects, or even animal and furniture, are too low on the first factor: they aren't specific enough to be useful. Very specific categories, like spaniel, sparrow, and armchair are too similar to other neighboring categories, like terrier, robin, and highchair. They are too low on the second factor. The idea that the basic level maximizes these two factors has been supported in several experiments by Murphy and Smith (1982). When people were taught hierarchies of categories that varied on only these two factors (with other factors held constant, which rarely happens in nature), they categorized objects fastest for the basic level names.

These two factors are important because of audience design. For the first factor this is clear: the more specifically an object is described, the easier it can be identified unambiguously. But if this is so, why isn't it always better to call something a terrier, sparrow, or armchair rather than a dog, bird, or chair? Why should the second factor make a difference? The reason is that using a specific name would imply, by Grice's (1975) maxim of quantity, that the distinction between terriers and other dogs, or between sparrows and other birds, or between armchairs and other types of chairs, was important in that situation. For Anne to tell Pierre, Sit in the desk chair, she must be distinguishing the desk chair from other types of chairs in the situation. If there is no other chair present, then her reference seems odd. Why is she emphasizing that it is not just any type of chair, but a desk chair? When a woman who owns only one car says, My Ferrari is out front, we take her as trying to impress us by distinguishing her Ferrari from more ordinary cars. Because basic level terms do not make such fine distinctions yet are still relatively informative, they are the neutral description.

But as Murphy and Brownell (unpublished) have argued, when the referent is atypical of the category to which it belongs, it generally needs a more specific description. People should refer to a penguin as the penguin and not as the bird, even though bird is the basic level term. Why? Penguins are not very similar to other birds, and so distinguishing penguins from other birds will be important for identifying them in most situations. To refer to a penguin as the bird might imply that it could fly, sing, and build nests, as most typical birds do. Indeed, people in Murphy and Brownell's experiments were able to categorize atypical objects faster with more specific names (like penguin) than with basic level names (like bird), even though for typical objects they were fastest for the basic level names. These results show that a preferred level of naming is not simply a result of greater familiarity or informativeness of names. Speakers design their descriptions to convey as much information as possible, yet not to make unnecessary or inappropriate distinctions in that situation.

ANAPHORA

Although speech generally sounds complete, it isn't. Much is left out in both the logic and the syntax. Anaphora is the name for the set of phenomena in which linguistic elements are missing, sometimes with other elements substituted for the missing ones (see Hankamer & Sag, 1976; Sag & Hankamer, 1981). The best known type of anaphora is pronominalization, in which a pronoun (like he) is used instead of a fuller description of the referent (like the handsome gentleman over there). Closely related to pronominalization is nominal anaphora, as in this example:

- (4) The blue racing car crashed into a wall.
The vehicle was destroyed.

In 4, a second definite description (the vehicle) refers to the same object as an earlier, longer definite description (the blue racing car). Since it could have been used in 4 just as well as the vehicle, pronominalization and nominal anaphora are probably understood by similar processes. The principles we discussed under definite reference also apply to pronominalization and nominal anaphora when they are definite.

In understanding nominal and pronominal anaphora, listeners rely directly on the design assumption in drawing what Clark and Haviland (1977; Haviland and Clark, 1974; Clark, 1977a, b) have called bridging inferences. Consider this two-sentence segment of discourse:

- (5) Marie took some picnic supplies out of the car. The beer was warm.

The problem here is that the speaker refers to some beer in the second sentence even though there is none in the discourse model so far. The listener, however, assumes that there ought to be some beer there if the speaker designed the utterance so that he could find such a referent uniquely. So he adds some beer to the discourse model via the following bridging inference:

- (5') There must be some beer among the picnic supplies, and this is the beer referred to.

Similarly, consider these segments:

- (6) Anne: Look at that fellow dance!
Pierre: Yes, the mayor is an excellent dancer.

- (7) Jean had a bruise on his cheek. Marie did it.

For the mayor in 6 and for did it in 7, listeners would make the following

bridging inferences:

- (6') The fellow Anne referred to must be a mayor; he must be the mayor of the city Anne and Pierre are in at the moment; and he is the referent of the mayor.
- (7') Jean must have got a bruise on his cheek because someone did something to him, and what that someone did is the referent of did it.

The listener knows what bridging inference to add only by appealing to the design assumption (see Clark, 1977a,b). In 5 Pierre reasons this way: "If Anne had good reason to think that I could infer the presence of some beer uniquely, she must have thought the inference is the most obvious one available from our common ground. The beer cannot be in Marie's handbag or near the car or in the branch of a tree. It must be among the picnic supplies, the mutually obvious place for the beer to be."

There is good evidence that listeners do make such bridging inferences (see Clark and Haviland, 1977). First, they say they do. Second, they take time in drawing bridging inferences. Haviland and Clark (1974) compared the time people took in trying to understand the second sentences in 5 and 8:

- (8) Marie took some beer out of the car. The beer was warm.

In 5, listeners have to draw the bridging inferences in 5', but in 8, they need to infer only 8':

- (8') The beer mentioned in the first sentence is the beer being referred to in the second.

The more complicated bridge in 5' took about a fifth of a second longer to construct than the simpler bridge in 8' (see also Clark and Sengul, 1979; Carpenter and Just, 1977; Just and Carpenter, 1978; Lesgold et al., 1979). And finally, listeners often misremember having heard a bridging inference when it didn't in fact occur (Kintsch, 1974, Chapter 8).

Audience design also plays a role in understanding more complicated types of anaphora, as we can illustrate with the form called "sluicing" in 9:

- (9) Anne: Claude climbed the old clock tower and then waited for someone to rescue him.
 Pierre: I wonder why ____.

Examples like this have generally been analyzed as cases of deletion of a verb phrase or clause through a simple syntactic mechanism (Sag & Hankamer, 1981). In this system, Anne would interpret Pierre's blank by consulting the previous utterance for a syntactically suitable replacement, here probably the whole previous sentence, namely "Claude climbed the old clock tower and then waited for someone to rescue him." But suppose that she and Pierre mutually knew that Claude climbed up and down the old clock tower once a week. Now she would assume instead that the blank referred to only her second clause--"I wonder why Claude waited for someone to rescue him." Or suppose that they mutually knew that Claude was a parachutist who liked to drop onto tall buildings and wait to be rescued. This time she would assume that the blank referred to the first clause alone--"I wonder why Claude CLIMBED the old clock tower." The interpretation of anaphora of this and related types, then, also demands that listeners reason from their common ground about what the speaker could have expected the addressees to interpret uniquely.

So, audience design is crucial to many, and perhaps all, types of anaphora.

Without it, listeners couldn't resolve ambiguities or build the bridging inferences required by many types of anaphora. Syntactic information is almost never enough.

WORD MEANINGS

In many standard views of meaning (e.g., Katz, 1972), a word like bank appears in the lexicon with one or more lexical entries associated with it. Each entry contains, among other things, a list of the possible senses with which that word can be used. Bank would have two lexical entries, one for the senses clustering around "steep natural incline" and another for the senses around "financial institution," such as "financial institution," "building housing a financial institution," and "the reserve chips in certain betting games." Bank is considered a homonym since its two lexical entries are thought to be independent of each other.

This view of meaning is the one that psychologists have generally adopted in their models of word understanding. In most discussions of lexical access (Forster, 1979; Garrett, 1978; Swinney, 1979), for example, listeners are assumed to access, or activate, more than one sense of each word they hear and then select among these senses on the basis of context. Metaphorically, listeners thumb through their mental lexicons for the right lexical entry and then try out each listed meaning for its plausibility in context, eventually selecting the right one. Precisely how this is done differs from model to model, but otherwise there is considerable agreement. Most of the findings in lexical access have been interpreted from this point of view.

Audience design would play a role in these models, but its role would be limited. Speakers would have a particular sense of a word in mind, like "building housing a financial institution" for bank, and they would provide enough context to enable their addressees to select the intended sense rather than some other. Ordinarily, this would be easy, since there are only a few possible senses for each word, and the intended one is constrained by many factors at once.

But audience design is more central to word meaning than this. The vocabulary of a language is open ended, and novel coinages, or word innovations, can be added at any time. Imagine Anne saying to Pierre:

(10) Was your newspaper properly porched this morning?

Although porch exists in English as a noun, it doesn't exist as a verb. But if it doesn't, how can Pierre interpret it? He cannot consult his mental lexicon, which has no lexical entry for this novel coinage. He is forced to create a meaning for the word on the spot. The same problem arises in this small sample of innovative expressions (underlined):

(11) My sister managed to Houdini her way into her office this morning.

(12) The police car sireded up to the accident.

(13) I postcarded all my friends that I was having a good time in Aix.

As Clark and Clark (1979) argued, certain types of novel coinages are not merely ambiguous. They have, in principle, an indefinitely large number of possible senses. In different situations, to porch can be taken to mean "turn into a porch," "supply with a porch," "disarm with a drawing of a porch," and so on indefinitely. Its possible senses are not enumerable. There is no algorithm that can enumerate just those senses it could have and no other senses. Clark and Clark called words whose senses are not enumerable contextual expressions.

Contextual expressions are interpretable only because of their audience design. As Clark and Clark argued, a speaker uses such an expression only when he has good reason to believe that his addressees can readily identify the intended denotation uniquely on the basis of their common ground. When Anne uses to Houdini, she is confident that Houdini and his ability to go through locked doors are common ground to her and Pierre and that he can use this knowledge to come quickly and uniquely to the intended denotation "go through a locked door as if by magic." In many situations, the design assumption would be crucial to Pierre's reasoning. Without it, he may be able to identify many possible interpretations, yet not decide which one had been intended.

If contextual expressions were on the periphery of language, they would be little cause for concern. But they are ubiquitous, and they are understood quickly and readily in the normal course of conversations. There are many types of contextual expressions, including novel coinages in these constructions (Clark, 1982).

indirect nouns, as in I saw a Henry Moore today
 noun compounds, as in I just bought a bed clock
 genitives, as in My bus today was crowded
 denominal nouns, as in He is a Churchillite
 denominal verbs, as in The newsboy porched the newspaper
 eponymous verbs, as in I did a Napoleon for the camera
 pro-act verbs, as in Alice did the lawn
 denominal adjectives, as in He's very Churchillian
 non-predicating adjectives, as in This atomic book is interesting
 eponymous adjectives, as in She's very San Francisco

Contextual expressions are an ordinary part of language (see Clark & Clark, 1979; Clark, 1982; Nunberg, 1979). They are usually understood quickly and without special effort or disruption. In experiments by Gerrig and Clark, novel denominal verbs took only a fraction of a second longer, if that, to understand than analogous well established words. Most of the time, listeners don't even notice they are interpreting novel coinages, and when they do, they don't find them odd or unnatural. For a theory of understanding to be correct, therefore, it must deal with contextual expressions in the normal course of processing.

Current theories of understanding, however, don't (see Clark, 1982). All of them assume, in one way or another, that listeners select among senses already in the mental lexicon or, failing that, among senses that can be created by a "lexical rule." Neither of these solutions can cope with the indefinitely large number of senses of contextual expressions nor with the fact that these senses cannot be constructed by rule.

If these arguments are correct, audience design should play an essential role in the process of understanding words, and it does. In several experiments (Clark, unpublished), people were asked to interpret novel expressions such as these:

- (14) The photographer asked the man to do a Napoleon for the camera.
 (15) I have always wanted to do a Nancy King to the salad.

To interpret these, people have to know who Napoleon and Nancy King are, to know acts they have done, and specifically to know some act that fits these utterances uniquely--an act that the speaker would have good reason to think they could identify readily and uniquely. For cases like 14, this process can succeed, and most people in the experiment interpreted do a Napoleon as "tuck the hand inside the coat à la Napoleon." For cases like 15, the process can only fail--unless people are told something about Nancy King, which is a

made-up name not part of common ground. When people in this experiment were told a story about Nancy King from which they could create a coherent unique interpretation, they accepted 15 and interpreted it as predicted. When they were told a story that precluded such an interpretation, they rejected the utterance and gave no interpretation. So people made direct use of the design assumption in deciding what these words meant.

Since novel coinages depend crucially on audience design for their interpretation, they should be understood quickly when the context is appropriate and not so quickly when it isn't. Of course, an appropriate context should speed the processing of any word, but it should be even more helpful for novel coinages. To test this prediction, Gerrig and Clark (unpublished) had people read sentences such as the following:

(16) She had a special way of spicing omelettes.

(17) She had a special way of herbing omelettes.

Although both spice and herb are verbs created from nouns, the verb spice is already well established in the lexicon, while the verb herb is novel. Consider what should happen when 16 or 17 is preceded by 18 or 19:

(18) Eve was very good at using the right seasoning in her cooking.

(19) Eve was very good at cooking a variety of dishes.

With 18 but not 19, the speaker mentions seasoning, which makes the meaning of the verb herb easy to create. The mention of seasoning, however, should give less help to the verb spice, whose meaning doesn't need to be created in the same way: its meaning is already in the mental lexicon. These predictions were confirmed in the reading time of the two types of sentences.

Listeners can never know, when they hear a word, whether it is being used in one of its well established senses, if it has one, or in an innovative sense (see Clark, 1982). They may be certain that words like Houdini and hammer-pocket are innovative and require the creation of a new sense, since these are not in the mental lexicon. They cannot be certain for the rest. Ultimately, they must always rely on the design assumption in deciding what the speaker meant.

CONCLUSIONS

We have argued as follows. When people talk, they tailor what they say to their audience. They expect their audience to use that fact in figuring out what they meant. Since listeners can assume that the speaker had good reason to believe that the addressees could fully understand what he meant on the basis of their common ground--this is the design assumption--they can reason through to the speaker's meaning accordingly. We have illustrated this form of reasoning for definite reference, anaphora, and word innovations. Audience design, therefore, must play an essential role in theories of understanding.

But the models of understanding most prominent today make little or no use of audience design. They work primarily from perceptual data--the speech sounds--applying various strategies to identify phonetic sequences, words, syntactic constructions, and sentence meanings. When context is brought in, it is brought in to arbitrate among the possible interpretations created by the strategies working from the perceptual data.

This will not do. People have to reason from the design assumption even to

get syntax and word meanings right. For example, the syntax of to porch depends on which of the indefinitely large number of possible meanings it is intended to have on this occasion: Does it take a direct or an indirect object, an animate or inanimate subject, or any oblique objects? Listeners cannot decide without reference to audience design. Further, there is evidence in the Gerrig and Clark experiments that people trying to understand an utterance make use of the design assumption in parallel with its syntactic and semantic properties (see also Marslen-Wilson & Tyler, 1980). Audience design comes into play any time listeners make choices, and these occur at all levels of language, from the phonetic level up.

Why has audience design played so little part in psychological models of understanding? It is probably because there has been so much research on understanding of isolated sentences and so little research on conversations and other genuine communication. In research on isolated sentences, the beliefs and background of particular listeners are almost impossible to study. For psychological models to become truly psychological, they must bring in the thoughts of individual speakers and listeners. They must be more than models of language use. They must be models of language users.

ACKNOWLEDGEMENT

The preparation of this paper was supported in part by grant MH-20021 from the National Institute of Mental Health. We thank Eve V. Clark for her comments on earlier drafts of the paper.

REFERENCES

- [1] Carpenter, P. A., & Just, M. A., Integrative processes in comprehension, in: LaBerge, D. & Samuels, S. J. (eds.), *Basic processes in reading: Perception and comprehension* (Hillsdale, N.J.: Erlbaum, 1977).
- [2] Carroll, J. M., & Bever, T. C., Sentence comprehension: A study in the relation of knowledge to perception, in: Carterette, E. C. & Friedman, M. P. (eds.), *The handbook of perception*, Vol. 5. *Language and speech* (New York: Academic Press, 1975).
- [3] Clark, E. V., and Clark, H. H., When nouns surface as verbs. *Language*, 55 (1979) 767-811.
- [4] Clark, H. H., Bridging, in: Johnson-Laird, P. N. & Wason, P. C. (eds.) *Thinking* (Cambridge: Cambridge University Press, 1977a).
- [5] Clark, H. H., Inferences in comprehension, in: LaBerge, D. and Samuels, S. J. (eds.), *Basic processes in reading: Perception and comprehension* (Hillsdale, N.J.: Erlbaum, 1977b).
- [6] Clark, H. H., Making sense of nonce sense, in: Flores d'Arcais, G. B. & Jarvella, R. (eds.), *The process of language understanding* (New York: Wiley, 1982).
- [7] Clark, H. H., & Carlson, T. B., Context for comprehension, in: Long, J. and Baddeley, A. (eds.), *Attention and performance IX* (Hillsdale, N.J.: Erlbaum, 1981).
- [8] Clark, H. H., & Carlson, T. B., Hearers and speech acts. *Language*, 58 (1982).
- [9] Clark, H. H., & Carlson, T. B., Speech acts and hearers' beliefs, in: Smith, N. V. (ed.), *Mutual knowledge* (London: Academic Press, 1982).
- [10] Clark, H. H., & Haviland, S. E., Comprehension and the given-new contract, in: Freedle, R. O. (ed.), *Discourse production and comprehension* (Hillsdale, N.J.: Erlbaum, 1977).
- [11] Clark, H. H., & Marshall, C. R., Definite reference and mutual knowledge, in: Joshi, A. K., Webber, B. and Sag, I. A. (eds.), *Elements of discourse understanding* (Cambridge: Cambridge University Press, 1981).
- [12] Clark, H. H., & Sengul, C. J., In search of referents for nouns and pronouns. *Memory and Cognition*, 7 (1979) 35-41.
- [13] Cruse, D. A., The pragmatics of lexical specificity. *Journal of Linguistics*, 13 (1977) 153-164.
- [14] Forster, K., Levels of processing and the structure of the language processor, in: Cooper, W. E. and Walker, E. C. T. (eds.), *Sentence processing: Psycholinguistic studies presented to Merrill Garrett* (Hillsdale, N.J.: Erlbaum, 1979).
- [15] Garrett, M. E., Word and sentence perception, in: Held, R., Leibowitz, H. W. & Teuber, H-L. (eds.), *Handbook of sensory physiology*, Vol. VIII, *Perception* (Berlin: Springer Verlag, 1978).
- [16] Grice, H. P., Logic and conversation, in: Cole, P. & Morgan, J. L.

- (eds.), *Syntax and semantics 3: Speech acts* (New York: Academic Press, 1975).
- [17] Hankamer, J., & Sag, I. A., Deep and surface anaphora. *Linguistic Inquiry*, 7 (1976) 391-426.
- [18] Haviland, S. E., & Clark, H. H., What's new? Acquiring new information as a process in comprehension. *Journal of Verbal Learning and Verbal Behavior*, 13 (1974) 512-521.
- [19] Hawkins, J. A., *Definiteness and indefiniteness: A study in reference and grammaticality prediction* (London: Croom Helm, 1978).
- [20] Just, M.A., & Carpenter, P.A., Inference process during reading: Reflections from eye fixations, in: Senders, J. W., Fisher, D. F. & Monty, R. A. (eds.), *Eye movements and the higher psychological functions* (Hillsdale, N.J.: Erlbaum, 1978).
- [21] Katz, J. J., *Semantic theory* (New York: Harper and Row, 1972).
- [22] Kempson, R. M., *Presupposition and the delimitation of semantics* (Cambridge: Cambridge University Press, 1975).
- [23] Kintsch, W., *The representation of meaning in memory* (Hillsdale, N.J.: Erlbaum, 1974).
- [24] Krauss, R. M., & Weinheimer, S., Concurrent feedback, confirmation, and the encoding of referents in verbal communication. *Journal of Personality and Social Psychology*, 4 (1966) 343-346.
- [25] Krauss, R. M., & Weinheimer, S., Effect of referent similarity and communication mode on verbal encoding. *Journal of Verbal Learning and Verbal Behavior*, 6 (1967) 359-363.
- [26] Lesgold, A. M., Roth, S. F., & Curtis, M. E., Foregrounding effects in discourse comprehension. *Journal of Verbal Learning and Verbal Behavior*, 18 (1979) 291-308.
- [27] Marslen-Wilson, W., & Tyler, L. K., The temporal structure of spoken language understanding. *Cognition*, 8 (1980) 1-71.
- [28] Murphy, G. L., & Brownell, H. H., Typicality constraints on basic level superiority. Unpublished manuscript.
- [29] Murphy, G. L., & Smith, E. E., Basic level superiority in picture categorization. *Journal of Verbal Learning and Verbal Behavior*, 21 (1982) 1-20.
- [30] Nunberg, G., The non-uniqueness of semantic solutions: Polysemy. *Linguistics and Philosophy*, 3 (1979) 143-184.
- [31] Olson, D. R., Language and thought: Aspects of a cognitive theory of semantics. *Psychological Review*, 77 (1970) 257-273.
- [32] Rosch, E., Principles of categorization, in: Rosch, E. & Lloyd, B. B. (eds.), *Categorization and cognition* (Hillsdale, N. J.: Erlbaum, 1978).
- [33] Rosch, E., Mervis, C. B., Gray, W. D., Johnson, D. M., & Boyes-Braem, P., Basic objects in natural categories. *Cognitive Psychology*, 8 (1976) 382-439.

- [34] Sag, I. A., & Hankamer, J., Toward a theory of anaphoric processing, in: Stanford Working Papers in Cognitive Science (Stanford University, 1981).
- [35] Strawson, R. E., On referring. *Mind*, 59 (1950) 320-344.
- [36] Swinney, D. A., Lexical access during sentence comprehension: (Re)Consideration of context effects. *Journal of Verbal Learning and Verbal Behavior*, 18 (1979) 645-659.
- [37] Wilson, D., *Presuppositions and non-truth-conditional semantics* (London: Academic Press, 1975).

