A Framework for Comparing Language Experiences
(with particular emphasis on:
The Effect of Audience on Discourse Models)

Andee Rubin
Bolt Beranek and Newman Inc.
Cambridge, Massachusetts

Those of us who have been involved in the study of language have often thrown up our hands in dismay at the complexity of the problem (at least I have, almost daily) and tried somewhat desperately to find some facet of the many-faced gem we confront which appears manageable. This desire to focus - to train the flashlights we use to illuminate the problem on a well-circumscribed area - led, for example, to the early transformational grammarians' separation of syntax from the rest of language. Many researchers have since discarded that particular focus and attempted to integrate syntax, semantics and pragmatics in a single theory. Yet, even they end up focusing on a particular kind of linguistic interaction. Many study oral conversations, some look at computer-person dialogues, some study newspaper articles. These limitations in scope are often not explicit, but are reflected in the examples they choose and discuss. Very few people study more than one communicative situation and even if they do, they do not usually analyze the similarities and differences among them.

Just as the early transformational focus on syntax resulted in a model which missed many crucial insights about language, so does our current research risk formulating incomplete and even inaccurate models by focusing on certain communicative situations without adequate insight into their relationship to others. My own focus in an attempt to point out what such a narrow view might miss, and to provide a framework in which to examine a wide variety of language experiences and discover what effect the differences among them have on theories of language. This paper first presents the framework, then examines further one dimension of language experience -- the audience for an utterance -- as an example of the kind of considerations such an analysis suggests.

1. Building the Space of Language Experiences

Consider first two language experiences commonly studied by present-day investigators: face-to-face oral conversations and computer-person dialogues. These two situations differ in at least two ways: the modality of the interaction (current computer-person dialogues are written) and the lack of possibility of communicating with extra-linguistic devices such as gestures and facial expressions. The kinds of questions we want to be able to address are:

What are the effects of these two distinctions on the language used in each of these situations? What are the effects on the models of language use which we thereby formulate?

In order to capture these kinds of differences in a way which will enable us to approach these questions, I have used the metaphor of a multi-dimensional space. Each language experience lies at a point in the space defined by its position along several dimensions of the linguistic medium. The medium of a language experience is defined in contrast to its message; in as much as they can be separated, the message is what is communicated, while the medium is how it is communicated. Further, the medium here is expressed in experiential terms and does not represent simply the vehicle for the message; for example, the contrast is made between being in a conversation and watching a play, rather than between a conversation and a play. (1)

Consider as a starting point in building the space the following message: an invitation and directions to a party. One common linguistic situation in which this might be communicated is a face-to-face oral conversation. Conversations, however, do not need to be oral; it is quite possible to maintain all the aspects of a conversation while writing it down by, for example, passing notes. These two language experiences form a "minimal pair"; that is, they differ along

_____

(1) I have omitted from this discussion any consideration of the message communicated in a language experience. In Rubin (1978) I identify three message-related dimensions - structure, function and topic - and discuss their interactions with the medium-related dimensions introduced here.

only one dimension. We might represent the situation graphically by labelling the line connecting the two experiences with the dimension along which they differ. The relevant dimension in this diagram is modality i.e. whether a language experience is oral or written. (Modality will be further dissected below.)

| HAVING A GARDEN-VARIETY CONVERSATION	MODALITY	HAVING A CONVERSATION BY PASSING NOTES

Figure 1.

Let us now look at two other pairs of language experiences which illustrate another dimension. Consider communicating the same message over the telephone compared to using a tape cassette. In the first case, it is possible for the two participants to interact, for the listener to express confusion and ask for additional information, for the speaker to monitor the listener's reactions and provide a more complete explanation. In the cassette situation, the speaker must decide once and for all how to give the directions without the benefit of intermediate feedback; any feedback which might occur would happen after the listener had heard the tape all the way through and would be temporally removed from the time the speaker composed the tape. I have termed this dimension of language experience interaction, as Figure 2 illustrates.

The other minimal pair in Figure 2, which also illustrates the interaction dimension, is communicating by letter versus communicating by a conversation over teletypes (while this is a somewhat unusual communicative setting, many of the people reading this paper have probably participated in it). Here again the crucial difference between the two is the possibility of feedback. In this particular task, for example, the speaker might want to ask of the listener, "Do you know the corner of Lewis and Fairview?" and base her further explanation on the response. Such an exchange would be impossible in the case of a letter.

TALKING ON THE TELEPHONE	INTERACTION	LISTENING TO A TAPE CASSETTE

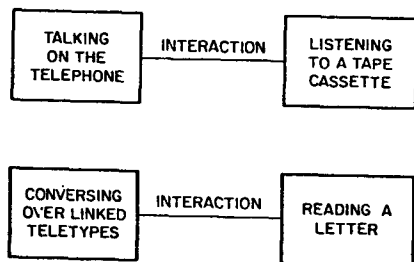| CONVERSING OVER LINKED TELETYPES	INTERACTION	READING A LETTER

Figure 2.

Notice now that we can connect the two minimal pairs in Figure 2 by lines labelled "modality". Reading a letter and listening to a cassette form a minimal pair which differ only in modality; the same is true for teletype and telephone conversations. The modality and interaction axes together form a plane in which we can place these four language experiences.
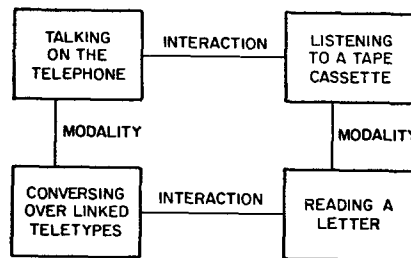

Figure 3.

Other dimensions can similarly be added to this space, by first discovering a minimal pair which focuses on a particular dimension, then attempting to pinpoint each language experience already in the space on that dimension and finally filling in the holes which exist because of the added axis. As an example of one step in building the space, consider the dimension of extra-linguistic communication, that is, communication by gestures, facial expressions, etc. For the message we are considering, gestures would be particularly useful to indicate spatial features such as "right" and "left" and the relative location of objects and landmarks. None of the four media in Figure 3 admit this type of interaction, but for each of them it is possible to construct an experience which differs from it along only this new dimension. For example, garden-variety conversations differ from telephone conversations because they allow this extra dimension, and passing notes differs from conversing over a teletype link in the same way. We now see where the pair of language experiences illustrated in Figure 1 comes in and by adding two more nodes we get the following cube.
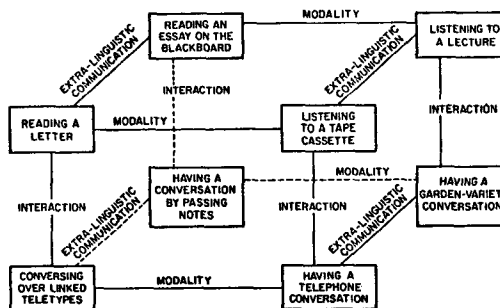

Figure 4.

The language experiences in this subspace differ in their degrees of naturalness; writing on the blackboard in such a way as to allow extra-linguistic communication but prohibit interaction, for example, seems contrived. This awkwardness is due primarily to the fact that certain dimensions generally covary and that the independence implied by a dimensional analysis doesn't really hold. Section 3 below discusses these interdependencies in more detail.

2. The Dimensions of Language Experience

In a similar fashion, we can add more dimensions to the space of language experiences. I have so far identified eight separate dimensions along which the medium of a language experience may vary. The dimensions are at least semi-independent; my informal criterion for listing a dimension separately was the existence of some minimal pair of language experiences whose media differed only along that dimension. The medium-related dimensions of language experience are: modality, interaction, extra-linguistic communication, spatial commonality, temporal commonality, concreteness of referents, separability of characters and specificity of audience. Below, each dimension is explicated by means of the question one would have to ask about a language experience to correctly place it on that dimension and additional details about its substructure and ramifications are given.

1. MODALITY - Is the message written or spoken? Even in this seemingly simple dimension are hidden at least two different characteristics which affect the communicative situation: prosody and permanence. I will briefly discuss these here, but a more extensive discussion of the components of modality may be found in Schallert, Kleiman & Rubin (1977) and Sticht, Beck, Hauke, Kleiman & James (1974) provide a review of the literature comparing auding (comprehension of oral language) and reading.

An obvious difference between an oral utterance and a text is the availability of prosodic cues. Temporal characteristics of speech such as pauses and changes in speed provide clues for the chunking of words into larger constituents. In general, pauses and breaths occur at syntactic boundaries and a more quickly spoken set of words may indicate an appositive phrase or an aside which is not germane to the top-level structure of the sentence. We rely on stress in oral language as an indicator of such discourse organizing topics as given/new, contrast and focus, as well as to aid in the disambiguation of pronominal references. Intonation is often used as an indication of the illocutionary force of an utterance or to communicate affective qualities of language such as humor or sarcasm.

While written text clearly lacks these properties, it has some compensating features. Punctuation and other textual devices provide a partial analogue of many prosodic features, including illocutionary force (. ? !), pauses (;), lists (, : ;), related statements (;) and contrast and emphasis (underlining and italicizing). A written message also provides the recipient with concretely indicated segments both on the lower levels of word and sentence and on the more abstract level of paragraph and section structure.

The second major distinction included in modality is the permanence of written text in contrast to the transitory nature of oral language. This permanence makes possible various "good reading" techniques such as skimming ahead to look at chapter headings, re-reading an entire paragraph whose point became clear only at the last sentence, or just re-reading a sentence which was misparsed the first time around. In oral language situations, such heuristics for dealing with misunderstanding are often replaced by an appeal to another (independent) characteristic of language experiences - interaction.

2. INTERACTION - Are the participants able to interact? In an interactive language experience, participants have the opportunity to indicate that they have not understood a previous utterance, that a pronominal or other reference is ambiguous or that they wish to change the topic. Keenan and Schieffelin (1976) in particular have represented the establishment of discourse topic as a dynamic process which includes input from all participants. This possibility for interaction means that there is less necessity for a participant to entertain and maintain a set of competing hypotheses about the meaning of some part of the message.

3. EXTRA-LINGUISTIC COMMUNICATION - Can the participants communicate via extra-linguistic means which require visual or tactile interaction? (This communication may sometimes be one-way, as in the case of a lecturer speaking to a large class.) Gestures, facial expressions and even body positions are all powerful in their communicative potential. In situations where emotions or spatial attributes are being communicated, these extra-linguistic means may be especially relevant. Children's early language experiences are especially dependent on this aspect of communication; deLaguna (1927) describes one developmental thread in children's language use as "a progressive freeing of speech from dependence on the perceived conditions under which it is uttered and heard, and from the behavior which accompanies it."

4. SPATIAL COMMONALITY - Can the participants interpret spatial deictic words such as "here" and "there" with reference to their own location? One indication that situations in which this condition is not met are difficult is the well-known situation in which two people, having arranged over the telephone to meet "here", discover that they had two different places in mind. Because the listener had to interpret the speaker's use of "here" relative to the speaker's location, it was necessary for her to know where the speaker was; incorrect information in this situation can have serious consequences. Young children may actually interpret "here" and "there" relative to their own position, rather than the speaker's. (see Tanz (1976) for details)

5. TEMPORAL COMMONALITY - Do the participants share a temporal context which allows for simple interpretation of temporal deictic terms such as "now", "today" and "last Sunday"? The correct interpretation of such words, as well as verb tense markers, requires the reader/listener to take the temporal point of view of the speaker/writer.

6. CONCRETENESS OF REFERENTS - Are the objects and events referred to visually present for the participants? If an object or event is concrete, many of its details are immediately apparent to the reader/listener besides the ones linguistically described in the message. Reading or hearing about an object which is not present often requires remembering a partial, incomplete description and then reformulating it as more information becomes available. Objects (or pictures) also provide an external "memory" for their existence and properties.

7. SEPARABILITY OF CHARACTERS - Are the distinctions among different people's statements and points of view clearly indicated? In face-to-face conversations, such distinctions are obvious, as each person makes his own statements and each point of view has a physical "anchor". In reading a play, characters' lines are clearly marked, although there is no physical object to which to attach each character. In a book, the reader must parcel out comments, feelings and motivations to characters on the basis of more subtle clues: punctuation, paragraph structure and inferences based on some consistent model of each of the characters.

8. SPECIFICITY OF AUDIENCE - How complete and specific is the speaker's model of the audience for her message? Two extremes which illustrate this dimension are garden- variety conversations in which the speaker and hearer know each other well and books, which are written for wide non-specific audiences. In

the former case, references to shared knowledge are possible, such as "The man looked like Uncle Joe," while in the latter, such an attempt would surely miss a large portion of the audience. To make matters worse, often a writer (or speaker) does not know who the audience is likely to be and in the case of books which are several hundred years old, the intended audience differs in significant ways from current readers.

Now that we have these eight dimensions, we can use them to generate new language experiences which begin to fill up the space. Watching a play, reading a book with pictures, viewing a movie with subtitles, reading a comic book - all these fit in the eight-dimensional space we have defined. (In Rubin (1978) I discuss quite a number of "intermediate" language experiences and show how they fit into a multi-dimensional space.) However, some areas of the space are only sparsely filled. These relatively empty sections are indications of interactions between dimensions; given a particular position on one dimension, the choices for certain other dimensions may be sharply constrained. Descriptions of some of these interactions follow.

3. Interactions Among the Dimensions

One fairly obvious interdependency is between spatial commonality and extra-linguistic communication. Since both rely primarily on the participants being in the same place, it is not surprising that most language experiences which exhibit the potential for extra-linguistic communication also allow participants straight-forward use of spatial deictic terms. (In fact, in Rubin (1978), I treat the two as a single dimension.) However, in a note left, for example, on the kitchen table, the writer may use "here" and "there", but cannot use gestures or facial expressions, so the two dimensions do not always co-occur.

Extra-linguistic communication is also most commonly found in oral language situations. The situation with the most potential for combining extra-linguistic communication with a written message is that of two people passing notes in what amounts to a written conversation. Although it is theoretically possible for them to point and grimace, it would be difficult for them to coordinate these gestures with the words in the written text.

Interaction and temporal commonality also appear closely linked. If the participants are not communicating in "real time" - that is, if the sending and receiving take place at different times - then it might even seen impossible for them to interact. However, if we allow the kind of attenuated interaction that takes place in, for example, an exchange of letters, we can maintain these dimensions as at least semi-independent.

Finally, we note that specificity of audience and interaction have an interesting relationship. Less well-defined audiences tend to occur in situations in which interaction is difficult, if not impossible. In lectures, for example, the speaker has only a vague idea of the audience's beliefs and interaction between them is limited. This covariance reflects two different facts. One is that in a large (and therefore poorly-specified) audience, interaction is restricted simply because of its size. The other is that interaction is one device by which speakers construct better models of their audiences; thus, a lack of interaction would lead to less well-specified audiences.

An obvious question at this point is: Why bother to separate dimensions that are so closely related? There are really two answers: the first methodological, the second historical. In terms of getting a clean model of the complex tangle of language experiences, it is better to postulate a large number of dimensions and specify how they interact than to identify only a small number but talk about subdimensions. Having a larger number of dimensions also inspires a wider range of language experiences when the process of filling in the space is carried out. Without the separation of temporal commonality and interaction, for example, we would have missed the subtle notion of temporally attenuated interaction.

The historical explanation derives from the original motivation for this work, which was an attempt to assess the relevance of children's early language experiences to their learning to read. Even if the dimensions identified here interact significantly, each still represents a cognitive skill which a child must learn in making the transition from garden-variety conversations to reading a text. In this framework, interactions among the dimensions are interesting because they represent pairs of skills which the child may have to learn together, rather than being able to separate them and learn one at a time.

Now that we have a notion where garden-variety dialogues fit into the framework of language experiences, it is possible to see what kinds of considerations we are liable to leave out if we focus only on conversations. While all of the dimensions identified above point out areas which deserve attention, I want to focus here on specificity of audience as an example of the ways our models must be stretched to account for the diversity of language experience.

4. Limitations and Compensation in Language Experiences: Non-Specific Audiences

Certain language experiences present problems for the participants, especially in comparison with garden-variety conversations, which have many communication-facilitating features. Lack of spatial commonality, for example, poses extra difficulty in the interpretation of certain deitic words, and the absence of non-verbal communication in telephone conversations makes expressing emotion especially hard. In some cases, an aspect of the medium itself provides compensation for the limitations. Written text, for example, partially compensates for its lack of prosodic cues to structure by its permanence, which allows the reader to make several attempts at parsing and understanding the words on the page. In other cases, we ourselves take into account the limitations of the medium by expressing our message differently. In talking on the phone, for example, we express our emotions more explicitly, rather than relying on facial expressions to communicate them in more subtle ways.

An important facilitating aspect of most garden-variety conversations is that they take place between people who have fairly good models of one another and who share a large set of beliefs. (See Cohen (1977) and Clark & Marshall (1978) for details on shared beliefs.) The disappearance of this feature in other language experiences can cause difficulties which require special attention from both speaker and listener. To get a feeling for the effect of an audience, consider the task of explaining the difference between analog and digital computers. Talking to a technically unsophisticated person makes the task hard enough, but it would be even more difficult if the audience were a large number of people with widely varying technical backgrounds. When one is faced with communicating with a person about whom one knows very little or, worse yet, an audience made up of many people with different beliefs it becomes necessary to use several compensating techniques to ensure that the message gets across. I will describe below some of the heuristics both speaker/writers and listener/readers use to deal with complex and poorlyspecified audiences. (From here on, I will use the words "speaker" and "listener" to refer to speaker/writers and listener/readers, respectively.)

5. Speakers' Heuristics for Complex Audiences

The audience for an utterance may be poorly specified in two different ways: it may be a single person about whom the speaker has less than complete knowledge or a group of people, each of whom the speaker knows more or less well. The speaker's task is to construct an utterance which is comprehensible to those who perceive that they are part of the intended audience. The following are techniques by which speakers may accomplish this task.

Identify the Audience: In some cases, the speaker really wants to address her remarks to a single person, even though several people are physically present and "available" as audience. Straightforward techniques exist for identifying the audience in these situations: a speaker may simply look in the direction of the intended audience or address

her by name. In giving a technical talk to a large and varied audience, a speaker may analogously select a subset of those present as the audience for a particular remark by using a phrase like "for you linguists in the audience". A more subtle and interesting method for accomplishing the same goal is to include in the remark a reference which only some of the audience understands, thereby clueing the others in to the fact that they may not get anything out of the utterance. At a recent conference attended by linguists, computer scientists and psychologists, a speaker, in answering a question from a computer scientist, resorted to some technical language (what he actually said was, "It's EQ, not EQUAL"), even though he knew two-thirds of the audience would be lost. Afterwards, he remarked that he had also realized that it was precisely the computer scientists who were confused about the point he was trying to make, so the remark was doubly appropriate: it selected exactly the audience who needed to comprehend it.

Play It Safe: Cover the Audience: If he realizes that the audience consists of two or more definable subgroups, a speaker may choose to include several descriptions of the same topic, one for each set of people. In addressing an audience made up of computer scientists and psychologists, a speaker might refer to the same concept by two different terms, e.g. "cache memory" for the computer scientists and "working memory" for the psychologists. In this case, most members of the audience understand only one of the two descriptions so both are necessary. A slightly different situation exists when a speaker makes a statement such as, "He had eyes just like Paul Newman...deep, dark blue." Here the elaboration may be seen as a comment to those in the audience who don't know what Paul Newman's eyes look like. Those who do must realize that, in some sense, the elaboration was not directed at them since it was planned for listeners with different knowledge. A third example of the "play it safe" strategy is this conversation which took place in front of an audience:

> B: Jerry has been studying the same thing.

> M: That's right...that's Jerry Fodor...I read his paper.

Had B and M been conversing in private, the explanation that the Jerry being referred to was Jerry Fodor would have been unnecessary. However, M was aware that some people in the audience might not be able to figure out who was being referred to, so he played it safe and made the reference clear.

Embrace Ambiguity: Sometimes, the speaker is aware that an utterance may be interpreted in two different ways, but decides that the ambiguity is acceptable or even desirable. An example from a personal letter: "The weather has been beautiful...perfect for riding a motorcycle in the country." What the writer had in mind here was a particular time she and the addressee had taken a motorcycle ride; she wanted to allude to that event. If, however, he didn't remember the ride, the sentence still communicated a coherent, if less specific, message. The writer actually considered both of these possibilities and decided that either reading of the sentence was acceptable.

The following sentence taken from an advertising brochure shows a different kind of ambiguity: "The cane seats of a Mad River canoe provide excellent ventilation and drainage." The ad will be read both by people who already know that Mad River canoes have cane seats and by those for whom this is a new fact. For the latter group, the word "cane" obviously conveys new information; for the former group, perhaps, it focuses on that aspect of the seat which is relevant to the properties discussed. (I use the word "focuses" here in an informal sense, but it is similar to its use in Grosz (1977) and Sidner (1978). Again, the ambiguity exists because of the non-specificity of the audience, but both readings are acceptable and, in fact, desirable.

Rely on Interaction: The standard way speakers check whether or not they are being understood and modify their utterances appropriately is by interacting with the hearer. In the Paul Newman example above, the speaker could have watched for signs of recognition (a smile or a nod) from the listener which would have made the elaboration unnecessary. In an interactive situation, the motorcycle example could have been followed by "Do you remember that time?" Unfortunately, language experiences in which the audience is unknown or unknowable, such as books and lectures, are often those in which interaction is difficult. The presence of both a specific audience and interaction is a "positive feedback" situation in which communication is greatly facilitated. The absence of both, however, necessitates the adoption of some of the heuristics mentioned here - and even then, communication may be impaired.

The heuristics identified above assume that the speaker takes the responsibility for the efficacy of the communication. Speakers in general, though, can assume that they share with the listener certain communicative principles of the type explicated by Grice (1975). They can similarly assume that listeners have certain heuristics for determining what assumptions the speaker is making about the audience and whether or not they as listeners fit those assumptions. Because they have this faith in their listeners, speakers sometimes just "broadcast" a remark, leaving it to the listener to decide who the intended audience is. Some of these listener's heuristics for interpreting broadcast utterances are described in the next section.

6. Hearers' Heuristics for "Broadcast" Utterances

Integral to a listener's understanding of any utterance is a model of the speaker's model of the hearer. Where there is some mismatch between this model and the listener himself, it may be difficult for him to figure out what the speaker really intended to communicate. In language experiences where the listener suspects that the speaker is broadcasting -- that is, not being careful about specifying an audience - he must "broadreceive", using some heuristics for deciding whether or not he should consider the remark to be addressed to him.

One general technique a listener might use is to compare the intended effect of the utterance with his current state. If he has already fulfilled the effect, he can consider the remark to be addressed elsewhere. Commands have clear intended effects, so it is relatively simple for a listener to use this heuristic with such speech acts. A member of a congregation who is already standing when the minister says "Please stand up" will understand that the utterance is not meant for him, rather than yelling out, "I already am!". Signs are another medium in which this broadcast behavior is apparent. A sign asking patients to PLEASE REGISTER WITH THE RECEPTIONIST is clearly meant only for those who haven't already done so; if patients didn't use this heuristic, they would find themselves in an infinite loop of registering. A somewhat more complex example is the familiar airport announcement: Extinguish all smoking materials and have your boarding pass ready." which selects two different subsets of the audience which hears it - those who are smoking and those who do not yet have their boarding passes ready.

Deciding on the intended effect of an utterance is no mean trick, of course; the speech act literature, (e.g., Austin (1962), Searle (1969)), makes this clear. One interesting example of the interaction of these considerations with those of audience is a subway sign proclaiming SMOKING IS DANGEROUS TO YOUR HEALTH. The intended effect may be seen either as getting smokers to give up their habit or as telling smokers and non-smokers alike that smoking is not healthful. In the first case, the intended audience is smaller than in the second. Although this more restricted audience is implied by the use of "you" on the sign, it seems plausible that the informational effect and, therefore, the larger audience is intended as well. A non-smoker's interpretation of the sign is then a complex process, involving an awareness of the two possible audiences for the message.

Finally, and most obviously, speakers assume that hearers will make use of pragmatic clues to determine whether they are part of the intended audience. A volleyball player yelling "I'll set" assumes that only the members of her team will attend to the remark; she doesn't need to preface it with a direct address to her teammates. Similarly, the same player calling "You hit" hopes that the pragmatic context is strong enough that the one person who is the intended audience can identify herself.

7. Summary

This admittedly brief discussion of the specificity of audiences is meant to illustrate the kinds of considerations a narrow focus on a single language experience might overlook. It is clear from just these few examples that the process of planning a speech act must utilize heuristics like those listed above and that speakers' models of listeners must contain some explicit representation of the size and specificity of the audience. These insights would not have arisen had we restricted ourselves to two-person conversations. The multi-dimensional space developed in this paper provides eight dimensions which can provoke similar investigations and a framework in which to integrate the results.

References

Austin, J.L. _How to do Things with Words_. Harvard University Press, Cambridge, Massachusetts, 1962.

Clark, Herbert H. and Marshall, Catherine. _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.

Cohen, Philip R. _On Knowing What to Say: Planning Speech Acts_. PhD Thesis, University of Toronto, Computer Science Department, 1978.

deLaguna, G. _Speech: Its function and development_. College Park, Maryland: McGrath Co., 1970 (Reprint of 1927 edition).

Grice, H.P., Logic and Conversation. In D. Davidson and G. Harman (Eds.) _The Logic of Grammar_, Encino, California: Dickenson Publishing Company, 1975.

Grosz, Barbara J. _The representation and use of focus in dialogue understanding_, PhD Thesis, University of California at Berkeley, Computer Science Department, May 1977.

Keenan, E., & Scheffelin, B. Topic as a discourse notion. In Li (Ed.), _Subject and topic_. Academic Press, 1976.

Rubin, Ann D. A Theoretical Taxonomy of the Differences Between Oral and Written Language. To appear in R. Spiro, B. Bruce and W. Brewer (Eds.), _Theoretical Issues in Reading Comprehension_, Hillsdale, N.J.: Lawrence Erlbaum, 1978, in press. Also as Center for the Study of Reading, Technical Report No. 35, 1978.

Schallert, D.L., Kleiman, G.M., & Rubin, A.D. Analyses of differences between written and oral language. Center for the Study of Reading, Technical Report No. 29, April 1977.

Searle, J.R. _Speech acts: An Essay in the philosophy of language_. Cambridge: Cambridge University Press, 1969.

Sidner, Candace L. _A Progress Report on the Discourse and Reference Components of PAL_. AI Memo 463, Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, Massachusetts, 1978.

Sticht, Thomas G., Beck, Lawrence J., Hauke, Robert N., Kleiman, Glenn M., and James, James H. _Auding and Reading: A Developmental Model_, Human Resources Research Organization, Alexandria, Virginia, 1974.

Tanz, C. _Studies in the acquisition of deictic terms_, PhD Thesis, University of Chicago, Department of Psychology, August, 1976.