# MANAGING THE FLOW OF IDEAS:

# A LOCAL MANAGEMENT APPROACH TO MESSAGE DESIGN

Barbara J. O'Keefe and Bruce L. Lambert

October, 1993

Barbara J. O'Keefe is an associate professor in the Department of Speech Communication, University of Illinois at Urbana-Champaign. Bruce L. Lambert is an assistant professor in the Department of Pharmacy Administration, University of Illinois at Chicago.

Author Address: 702 S. Wright St., #244, Urbana, IL 61801

Running Head: Local Management Approach

# MANAGING THE FLOW OF IDEAS:

# A LOCAL MANAGEMENT APPROACH TO MESSAGE DESIGN


## Abstract

As research in communication begins to develop serious hypotheses about the cognitive processes underlying communication, it will become increasingly important to be able to develop detailed models that can be methodologically realized in either experimental paradigms or computer simulations. This paper develops the case for a local management model of message design and presents an initial model of the message design process that offers an integrated treatment of message production, adaptation, and effects.



MANAGING THE FLOW OF IDEAS:
A LOCAL MANAGEMENT APPROACH TO MESSAGE DESIGN


Our objective is to develop a theory of message design, a systematic theory of the relationship between message structure and message function. Message structure refers to the substance, organization, and placement of discourse. Message function involves both the antecedent conditions of message generation (especially the goals of the message producer) and the intended and unintended effects of the message.

Most current theories of message design are based on a holistic, functional view of messages. In holistic-functional analysis, units in the stream of discourse are identified and then labelled, as wholes, in terms of their discourse function. In different holistic-functional approaches, the nature of the unit may vary (it may be a turn, it may be a text) and the kind of functional characterization may vary (illocutionary force, perlocutionary effect, type of adjacency pair, type of compliance-gaining strategy, etc.). Examples of holistic functional categories include illocutionary and perlocutionary acts (Searle, 1969); taxonomies of communicative strategies (e.g., Marwell & Schmitt, 1967a, 1967b); first and second pair parts (Schegloff, 1968); and so on. Once a functional terminology has been constructed, it can be used in research on the conditions that influence the type of message produced and the consequences of using different types of messages.

This approach, in which functional theory is based on functional categorizations of discourse, has been fruitful in a variety of very different research traditions, including research on the development, deployment, and effects of interpersonal message strategies (for reviews, see McLaughlin, 1984; O'Keefe & Delia, 1988; Seibold, Cantrill, & Meyers, 1985) and research on the organization of face-to-face interaction (for reviews, see McLaughlin, 1984; Roger & Bull, 1989; Taylor & Cameron, 1987). However, while this approach has been helpful in generating information about the general pattern, distribution, and effects of messages, it is limited in its ability to address a set of important, basic questions about message production, variation, and outcomes.

In this paper we describe an alternative to the standard holistic-functional approach to message analysis, one that treats message design as the local management of situated beliefs. Rather than treating messages as coherent instantiations of globally defined actions, this approach treats messages as collations of thoughts. Message design, then, is conceptualized as the local management of the flow of thought--both the management of own thoughts by the message producer and the management of the other's thoughts in the service of communicative goals. The question we address is how patterns of thoughts and ways of managing thought can give rise to message structure and message functions.

The first section of the paper describes current, problem-solving models of message generation and their limitations in accounting for message design. The second section of the paper describes a set of phenomena that present special difficulties for a problem-solving approach and suggest the value of a local management approach. The third section provides a general characterization of the local management approach. The fourth section presents a detailed model of message design as the management of thought, and summarizes our work to date developing models of message generation. The fifth section shows how this model can account for message adaptation and effects. The conclusion discusses the broader implications of this work for theories of discourse and communication.

## The Problem-Solving Model

The most popular current models of message generation, which we call "problem-solving models," are generally based on a holistic-functional view of messages. In this section, we discuss models based on AI conceptions of planning as problem-solving (Fikes & Nilsson, 1971; Miller, Galanter, & Pribram, 1960; Newell & Simon, 1972; Sacerdoti, 1974) and the role of functional categories in such models. This approach has stimulated a great deal of research and is exemplified in impressive computer models of Cohen and Perrault (1979) and Appelt (1982). In what follows, we describe the problem-solving approach and discuss its limitations.

In a problem-solving model, goals are presented to the system; subgoals are identified by comparing the current state to the goal state, and actions are chosen to eliminate existing differences between current and goal states (see Levelt, 1989). For example, if the goal state specifies "H believes that S believes p," and the current state specifies "H believes that S believes not p," then the task is to select an action that changes H's beliefs. The choice of an appropriate action is made possible by indexing actions by the differences they eliminate.

A problem-solving model will have a repertoire of actions at its disposal, each indexed by the type of difference it eliminates. Within different models, these action repertoires will be represented a bit differently. For purposes of discussion, assume they take the form of if-then rules: "IF the intention is to commit oneself to the truth of p, THEN assert p" (Levelt, 1989, p. 10, emphasis in the original). This is unobjectionable, if one grants that asserting p, by definition, has the desired effect on H's beliefs (Cohen & Perrault, 1979; Searle, 1969).

From the standpoint of a theory of message design, a key weakness of a problem-solving approach derives precisely from the fact that problems are represented and solved abstractly, i.e., in terms of types of situations and actions rather than in terms of specific situations and the field of thoughts that accompany them. A plan derived from an abstract problem-solving process will always be a sequence of act types. To move from a sequence of act types to a message, each act type must be instantiated as a particular utterance.

Hence, these models must somehow solve the problem of instantiation: they must show how a system moves from a high-level act category (e.g., request H to do A) to a string of produced sounds. A common solution is to posit a hierarchy of linguistic structures in which a structure at one level is instantiated from choices at the next level down. For example, in Appelt's (1987) model, the system is endowed with a repertoire of choices at each of a number of levels. Having selected a speech act to instantiate (e.g., an assertion), the next task is to select a sentence-type. At the level of sentence type, the system can choose between declarative, interrogative, and imperative sentence types. The choice is made possible by indexing each sentence type by the speech acts it instantiates. Thus, declaratives instantiate assertions, interrogatives instantiate requests, and imperatives instantiate commands. Needing to instantiate an assertion, the system's choice of a declarative sentence is simple.

Essentially the same approach is taken at levels below the sentence (see Garrett, 1975, and the critique of Garrett by Dell & Juliano, 1990). Once a sentence type is chosen, words must be chosen to fill the slots in the sentence frame. The system has a lexicon of words, and each word is indexed by its grammatical function (noun, verb, determiner, etc.) and by its meaning. Once words are chosen, morphemes and phonemes must be chosen to fill in the slots in the word's representation. Once phonemes are chosen, articulatory patterns must be chosen to produce the relevant sounds. At each level, there is a choice to be made, and the dilemma of choice is always resolved in the same way, by indexing choices in terms available at the highest adjacent level of abstraction.

These abstract representations with slots to be filled are called functional representations (Garrett, 1988), and we refer to this way of solving the instantiation problem as a functional indexing scheme. The use of functional indexing schemes is a definitive characteristic of standard, problem-solving models of message production.

However, there are difficulties with functional indexing schemes as they have been implemented in computer models of message production. The indexing that is required in order to enable rational choice among forms depends on there being a decontextualized relationship between form and function. But, to put it bluntly, decontextualized linguistic forms have no functional significance. It is a truism to say that the meaning of a form depends on the context of its use. At the level of discourse acts or message features, the evidence shows that the form-function relationship is mediated by reasoning from context-specific beliefs (Levinson, 1982; Tannen, 1993). At the level of words, the analogous phenomenon is known as polysemy (Green, 1989). Similar context dependencies are apparent at lower levels of abstraction as well (Levelt, 1989).

Notice that the instantiation problem arises, in the very beginning, from assuming that the materials of planning--communicative acts--are represented, selected, and ordered independently from some specific context. Not surprisingly, this problem has been recognized by researchers working within the problem-solving approach, and various solutions have been proposed. For example, Hovy (1990) offers a model that assumes a particular input representation rather than generating the input through problem solving. However, his approach, since it involves reasoning with abstract message features rather than concrete situated message contents, still faces the instantiation problem, which he refers to as "casting." Even though Hovy's model offers a richer array of functional indices (e.g., he includes rhetorical indices that permit choice among instantiations based on situational features), it nonetheless must ultimately rest on a view of form-function relations as decontextualized and fixed.

[New material here?]

## Speech and the Flow of Thought

In the preceding section we discussed the difficulties faced by holistic-functional models that result from the assumption of a decontextualized relationship between message forms and message functions. In this section we discuss a second problem facing holistic-functional approaches, namely, that message structure and function are not holistic, but rather reflect the grounding of messages in an ongoing stream of thought and action. In this section, we illustrate the grounding of speech in the flow of thought by considering two distinct classes of messages: (1) descriptions and (2) complex interpersonal tasks. We show how these phenomena point toward a very different image of message design, namely, a view of messages as the local management of thought.

### The Structure of Descriptions

A good deal of thought and research has been devoted to understanding descriptions, although what is meant by description can vary quite widely from one investigation to another. For example, one tradition examines relatively short characterizations of objects or states of affairs (e.g., "the dog with the brown spots") and analyzes their properties and use in acts of referring (e.g., Grosz, 1981). Other lines of research focus on extended speech acts in which description is used, not simply in the service of reference, but to offer information about a referent to a hearer. While all this research on descriptions clearly connects to important underlying theoretical issues, nonetheless the kinds of texts, discourse structures, and communicative functions under examination vary widely across this topic area.

However, one consistent theme throughout the study of descriptions is the close relationship between message structure and the substance and organization of knowledge in the topic domain. For example, Grosz (1981) studied instructions provided for the assembly of a mechanical device; to explain the structure of the instructions she developed a model of describing as the movement of focus of attention through a knowledge structure. A similar model was developed in very different research contexts by Chafe (1979), who studied descriptions of events, Levy (1983), who analyzed descriptions of student class schedules, and Sibun (1990, in press), who modelled descriptions of kitchens and family trees.

Examples of a highly detailed analyses of discourse structure as reflecting the movement of attention through a structure of knowledge can be found in studies of descriptions of spatial layouts. Beginning with Linde and Labov's (1975) study of apartment descriptions, considerable research has shown that the structure of such descriptions, whether they describe a room (e.g., Ullmer-Ehrich, 1982), a residence (e.g., Linde & Labov, 1975), or a route between points in a city (e.g., Wunderlich & Reinelt, 1982), generally exhibit a pattern in which hypothetical movement through a space provides the organizing principle for a linearized recounting of spatial information. So, for example, both descriptions of apartments and rooms commonly take the form of a "tour" in which the hypothetical gaze of the hearer is guided through the space and led to focus on key features of interest.

As Ullmer-Ehrich (1982) argues, one should not assume that explaining the structure of descriptions is therefore simply a matter of developing a model of the representation of spatial information, since there are three problems that must be solved in moving from a spatial representation to a verbal message. The first problem, the selection problem, arises from the fact that much more spatial information is stored than can or should be expressed. Hence, the speaker must select just those elements of the representation that are relevant for the purpose at hand. The second problem, the transformation problem, arises from the fact that in order to be expressed, elements of the representation must be temporalized, placed in a one-dimensional order. The third problem, the symbolization problem, arises from the fact that spatial information must be verbally formulated.

All three of these issues can be addressed within a theory of focus. As Groscz (1981) envisions it, "focusing is the active process, engaged in by the participants in a dialogue, of concentrating attention on, or highlighting, a subset of their shared reality" (p. 101). A focus involves not simply attention given to cognitive elements, but a perspective on those elements, a perspective that is implicit in the terms of a description. By specifying the perspective, a theory of focus contributes to an analysis of symbolization; by modelling the restriction on relevance that comes with focusing attention, a theory of focus contributes to an analysis of selection; and by explaining the movement of focus through a knowledge structure, a theory of focus contributes to an analysis of the linearization of knowledge as text.

In the case of descriptions, much of the substance and organization of messages can be seen as reflecting the substance and organization of knowledge in the topic domain. This knowledge is re-presented, selected, and ordered for expression as focus follows a route (e.g., a tour) through the spatial representation. Hence, descriptions originate as content which becomes focal as attention is directed by a specific communicative task; utterances result from expression of focal elements.

This general picture of message production is also exemplified in work by Kellermann and her associates (for a review, see Kellermann & Lim, 1989) on a very different problem: exchanges of information in initial interactions. They have shown how the problem of becoming acquainted leads individuals to move systematically through an agenda of knowledge exchanges; the talk produced in discussing a topic reflects a standard agenda of points of focus within the topic. Although their work is presented in a framework very different from ours, nonetheless this work provides yet another example of how talk is generated by the movement of focus through knowledge structures.

## Complex Interpersonal Tasks and Message Design

Descriptions is a domain in which performance is generally uniform across individuals--different people produce messages that are substantially identical in structure. By contrast, complex interpersonal situations (regulating, comforting, etc.) elicit highly variable performance from different individuals.

Because of this, performance on complex communication tasks provides a useful point of contrast for a local management theory of message design. It might be argued, for example, that descriptions are one of the domains (initial interactions being another) in which behavior is highly routinized (indeed, Linde and Labov, 1975, make this claim) and therefore speakers make no detectable strategic choices. With tasks that that elicit greater functional variation, high-level choices among distinct strategies might be more apparent.

But in this section we argue that responses to complex communication tasks offer yet another compelling example of the grounding of talk in the ongoing stream of thought. We have conducted several recent investigations in which we have addressed the question of whether messages consist of relatively distinct and unified "strategies" or relatively inchoate and fragmented collations of thoughts. To address this question, we required a different type of message analysis than had currently been employed in research on message structure and effects; the standard approach, holistic-functional classification, obviously begs the very question we wanted to answer.

Our method is based on the identification of relatively fine-grained discourse units. Adapting techniques used by Chafe (1979) in studies of oral discourse and by Hunt, Matsuhashi, and others in studies of written communication (for a review, see Hillocks, 1986), we segmented messages into thought units, which essentially correspond to independent clauses. These thought units are then grouped based on synonymy into categories that reflect the basic idea expressed in the thought, independent of specific wording. The criteria for synonymy are quite conservative, and generally require that alternative members of a category be similar except for grammatical transformation or substitution of synonymous terms or phrases.

For example, Lambert (1992) studied messages elicited by a hypothetical situation in which a friend repeatedly breaks dates with the speaker. He analyzed 320 messages, and found that their content could be characterized as subsets drawn from a list of 72 basic types of thought units (see Figure 1). In subsequent studies, we have further classified thought units into larger categories--for example, Saeki and O'Keefe (1993) and Lambert and Lee (1993) grouped types of thoughts into themes based on similarity in meaning and coocurrence.

Using these methods, we have observed that the expression of thought units is not strongly constrained by the functional "type" of a message. In Lambert's (1992) investigation, there were some tendencies for particular types of thoughts to cooccur, but in general the thoughts did not cluster strongly. Some kinds of thought units were very common, and appeared in up to a third of the messages; other thought units occurred much less commonly.

In their similar investigation, Saeki and O'Keefe (1993) studied messages elicited by a hypothetical situation in which a student must tell another student that he or she has been rejected for admission to an honor society. The situation was varied in terms of the relationship between speaker and hearer (friend vs. stranger) and the hearer's qualifications (well vs. poorly qualified). Saeki and O'Keefe analyzed 228 messages produced by American and Japanese students and found that the messages could be characterized as subsets drawn from 21 basic types of thought units. They further grouped the 21 thought units into eight content themes (see Figure 2). Their results showed that elaboration of each of the eight themes was influenced by a distinctive set of situation features and associated with a distinctive set of goals. They concluded that this pattern of findings suggests that rather than being composed of unitary and coherent strategies, instead messages are "comprised of many independent parts" (p. 28).

Lambert and Lee (1993) investigated messages produced by pharmacy students for a hypothetical patient compliance situation. They analyzed 85 messages and found that the messages could be characterized as subsets drawn from a set of 61 idea types. Using procedures similar to those of Saeki and O'Keefe (1993), idea types were grouped into 11 distinct content themes (see Figure 3). The degree of elaboration of each theme within a message was calculated by summing the number of idea types associated with a given theme in the message. The degree of elaboration of a theme was not in general associated with the elaboration of other themes; only seven of 55 correlations were significant, and even these were relatively small effects (i.e., $r < .30$).

In general, then, the distribution of content themes within messages is relatively unconstrained by the functional type of the message. Any theory of message design must account for the fact that messages with very different points and effects can nonetheless share a good deal of content and that messages with similar points or effects can be very different in content.

This property of messages arises from the fact that message effects are associated with specific message contents rather than messages taken as wholes. For example, Lambert and Lee (1993) found that only three of the 11 content themes played a significant role in influencing perceptions of effectiveness in meeting task and interpersonal goals. [Other study too?]

## Conclusion: The Flow of Thought and the Flow of Talk

In short, then, it appears that communication situations should be represented as organized fields of thoughts. Message structures arise as focus moves through the field of thoughts. Focus is driven by goals and guided by the route that the speaker formulates to move through the field. In the case of descriptions, a route is envisioned only in relationship to a particular spatial representation; similarly, in the case of refusals, compliance-gaining messages, regulative messages, and the like, a route is formulated in relationship to a particular field of thoughts. The resulting message, rather than being a functionally unified act, instead is a collation of thoughts, each of which may have distinctive consequences and effects.

The diversity of messages in complex situations arises from the fact that communicators can have different goals and different "routes" associated with those goals. When this is true, it will naturally lead to a diversity in focusing; where there is diversity in focusing, there will be diversity in the thoughts selected for expression and differences in the resulting message.

Whether a message domain is characterized by low or high functional diversity, unity or fragmentation, message structure reflects the intimate relationship between knowledge and expression. Specifically, message structure results from the movement of a focus of attention through a structure of knowledge. A local management model thus offers a particular view of message planning as involving the imposition of a focus or perspective on a representation of the situation.

One key advantage of a local management view of planning is that it faces no instantiation problem--the route is traced in the space to be negotiated, not a generic representation of spaces; the actions are specific thoughts to be uttered, not a generic representation of actions that might be undertaken. Since planning occurs with the materials provided within the context and not with decontextualized functional categories, the problem of connecting an abstract representation of action with expressions appropriate to the context simply disappears.

A second key advantage of a local management view is that it does not presume message coherence. A problem-solving model can only generate messages that conform to a plan and cohere around a set of goals. But real messages often contain functionally distinct or even dysfunctional content. A local management approach can easily accommodate this fact, since the field of thoughts may or may not be functional and coherent. To the extent that an individual has conflicting or dysfunctional thoughts, he or she may produce a relatively incoherent or fragmented message.

## Local Management and Message Design

In the two previous sections we highlighted the need for an approach to message design that avoids holistic-functional analysis and instead treats message structure as the movement of focus of attention through knowledge structures. We also argued for a view of planning as a process of local management rather than problem-solving. In this section we outline an approach to message design that meets these requirements.

First, we describe our general image of planning. Second, since any model of message generation is predicated on some particular image of cognitive architecture, we explain why our model assumes a connectionist architecture.

### Planning as Local Management

Local management conceptions of planning (Agre & Chapman, 1987, 1990) and language production (Sibun, 1991b) derive from an ethnomethodological and activity-theoretic understanding of everyday action (see, e.g., Garfinkel, 1967; Lave, 1988; Suchman, 1987). Agre and Chapman (1990) execute a classic "ethnomethodological inversion" when they claim that planning is based on communication rather than vice versa: "Our ability to make and use plans is built on our ability to use language during activities we share with others" (p. 25). Planning itself is made possible only by our abilities to improvise and interpret sequentially unfolding, situated activities. Since improvisational skills are what make planning possible, they themselves cannot be explained in terms of planning. Rather, an independent theory must be offered to explain what allows people to reason through each moment's action by a fresh reasoning-through of that moment's situation (Agre & Chapman, 1990, p. 21). While the models we propose differ to some extent from the kind of cognitive machinery Agre and Chapman suggest, our perspective, like theirs, reflects an effort to explain the cognitive substrate that supports improvisation during situated activity.

The model we propose also shares design principles and theoretical commitments with Sibun's (1991a, 1991b) Salix, a system for generating natural language descriptions of houses and families, embodies. Salix generates coherent text without using abstract structural representations and without reference to act types or global plans. Instead, text is generated by exploiting the existing structure of the content domain being described. Global planning and organizational processes are supplanted by strategies for choosing what to say next from among a small set of locally available alternatives. In a similar vein our model explains message design (generation and effects) without reference to act types or hierarchical plans. Instead, message structure is explained in terms of the organization of situated beliefs and the movement of focus through belief structures.

### The Organization of Knowledge

There is, at present, a controversy in cognitive science concerning the fundamental design of the architecture of cognition. Our model assumes a connectionist (paralled distributed processing or PDP) architecture. This section outlines the basis for this preference.

Since the 1950's, it has been widely agreed that the mind is a "physical symbol system." The physical symbol system hypothesis (henceforth PSSH), derived primarily from the work of Newell and Simon (Newell, 1980), asserts that cognition is the creation, destruction, manipulation and transformation of symbols and complex symbol structures.

A symbol is a physically realizable atomic (primitive) element with no internal structure, occupying a discrete location in memory. Symbolic structures are collections of symbols which stand in rule-governed, formally specified, syntactic relations to one another. PSS's consist of processes that act on symbols and symbol structures. A unique symbol designates each unique process. Symbols are interpreted by the system in that a process can act according to a symbol or execute or manipulate the symbol or process it designates. Symbol systems evolve over time as they move through different symbolic states or spaces. Intelligent action on this account involves heuristic selective search through symbolic structures. Aspects of a wide variety of human abilities, including problem solving, learning, language production and comprehension, and vision have been simulated using this method of "symbols and search." The approach is frequently, but not necessarily, associated with serial computing and the von Neumann computer architecture.

During the last decade or so, this view has been challenged by a community of researchers who view cognitive architecture as being constituted by parallel, distributed representations and processes (Hinton & Anderson, 1981; McClelland, Rumelhart and the PDP Research Group, 1986; Rumelhart, McClelland & the PDP Research Group, 1986; Smolensky, 1988). Instead of symbols and search, the primitives here are units and connections. On this view, the mind consists of a highly interconnected network of simple computing elements functioning in parallel.

In a PDP representation, each element or unit has a numeric state of activation. Units are connected to one another by modifiable, weighted connections. In most connectionist models, the effect of a unit i on a unit j is the product of i's activation with the numeric strength of the weight connecting i to j. Information is represented in the network as vectors of activation values across input and output units. The activation of a given unit is typically interpreted as the degree of presence or absence of some concept or (micro-) feature of the input or output. The inputs effect the outputs by propagating their activations through the network of weighted connections. The network's knowledge is contained in the matrix of weights that connects units to one another.

Several attractive computational properties emerge naturally from parallel distributed processing models. Among those frequently cited are (a) content-addressable memory, (b) graded, continuous processing, (c) context-sensitivity, and (d) learning (Anderson & Hinton, 1981; Clark, 1989; McClelland, Rumelhart, & Hinton, 1986; Mikkulainen, 1993).

Many of these properties result from the use of distributed, rather than local or discrete, representations of items in memory and from storing knowledge in connections between units rather than at discrete locations in memory. In a standard symbolic model, an item in memory is represented by an atomic symbolic token residing at a specific address in memory. In a PDP model, by contrast, an item is represented by a pattern of activation over many units. Where a standard model might represent a communicative goal as a discrete symbol, a PDP model can represent a goal as a pattern of activation over many units, where each unit represents the degree of presence or absence of a specific thought.

This mode of representation and processing allows items in memory to be content addressed. One can retrieve a memory simply by activating parts of the memory. The knowledge associated with an item is not stored at any discrete location, but rather in the connections between units that represent microfeatures of the item being remembered.

In standard, symbolic AI, a symbol is either present or absent in an all or nothing manner. Actions are either taken or not taken based on the presence or absence of discrete symbols. This all-or-nothing mode of response leads symbolic models to manifest brittle performance in the face of exceptional or novel input. However, continuous shadings of meaning and response are possible in a PDP framework. Units representing microfeatures of a stimulus object may take on a range of continuous values representing degrees of presence or absence. In addition, since meaningful objects (e.g., goals) are represented by patterns of activation over many units, slightly different goals may be represented by slightly different patterns of activation. Thus, graded, continuous processing is manifest at the level of individual units and at the level of patterns of activation.

An appealing form of context sensitivity also emerges from this type of graded, distributed processing. As Smolensky (1988) puts it, the context of a symbol in a standard cognitivist model is provided by other symbols; whereas, the context of a PDP "symbol" is part of the internal structure, the distributed pattern of activation, which is the representation of the symbol itself. Thus, the "same" symbol is represented in different contexts by slightly different patterns of activation across the subsymbols or microfeatures that comprise it.

Finally, PDP models are able to learn from examples and generalize their knowledge to novel inputs. There are procedures for adjusting the connection weights such that they come to approximate the function embodied in a set of correct input-output pairs. Most learning procedures follow the same general pattern (Hinton, 1990). First, a vector of input activations is presented to the input units, and activation is propagated forward to the output units. Next, the desired output value (from the correct example) is compared to the actual value produced by the training example, and the difference is computed. Finally, each of the weights is modified in proportion to the error it caused. The process is repeated until the weights reach a state which minimizes the square of the errors across all of the units and examples. A simple procedure with only input and output units has principled limitations (Minsky & Papert, 1969), but there are now learning rules for networks with hidden units that can approximate any input-output function.

In sum, several important computational properties emerge naturally out of PDP-style representations and processes, and these properties are difficult if not impossible to implement with standard symbolic AI methods. There are other attractive features of PDP models (such as automatic default assignment and spontaneous generalization, see Clark, 1989), but our choice of PDP over traditional models was motivated in large part by the properties just described.

## A Model of Message Generation

In this section, we propose a specific model of message generation as local management of situated beliefs. We describe the antecedents of the model in work by the Mannheim group on speech and situation; we sketch a view of relevance as a function of focus; and we discuss development of a computer simulation by Lambert (1992) that provides a concrete demonstration of a key process in our model.

### Antecedents of the Model

Our model is influenced by the Basic Model of Speech Production created by Herrmann (1983) and his colleagues at the Research Group in Language and Cognition at the University of Mannheim, Germany

(Hoppe-Graff, Herrmann, Winterhoff-Spurk, & Mangold, 1985). While Herrmann's (1983) model seems the most useful starting point for this project, it should be noted that many of the central features of the model are similar to other models reviewed in Levelt (1989). For the present purposes, the key feature of all these models is that they posit a selection process in which a subset of available knowledge about a topic is organized for expression during a given turn at talk.

Herrmann's (1983) model is a multi-step, cognitive model of language production in which message generation is seen as originating in the construction of a situation representation by a situation-interpreting program (see Figure 1). The situation representation contains both information about relevant situation parameters as well as procedural and declarative knowledge retrieved by the situation interpreting program. Situation parameters are features of the immediate context that will shape the choice of appropriate utterance (e.g., in a request situation: hearer's ability to perform the requested action, hearer's willingness to perform, speaker's legitimate right to make the request, etc.). The remaining knowledge retrieved from long-term memory by the interpretation program consists of situation-specific scripts and more loosely organized facts that embody knowledge about the constraints on and procedures appropriate to the current situation. The declarative portion of this accumulated store of knowledge is referred to as the propositional base (PB) of the utterance.

According to Herrmann (1983, p. 25), the propositional base of an utterance is "the foundation of what is meant . . .[including] all that the speaker has perceived, recollected, imagined, inferred, presumed, planned, etc. in connection with the process of speaking . . . [it is] the activated data base of the utterance." The PB consists of those "elements of knowledge" relevant to achieving the speaker's current goal. The PB comprises both what the speaker means (his/her intention) and what the hearer must reconstruct in order to comprehend a given utterance. The PB does not always come in neat, pre-formed chunks from memory, but may require active construction processes for its on-line elaboration.

If the PB is the foundation of what is meant, the semantic input (SI) is the foundation of what is said (Herrmann, 1983). According to Hoppe-Graff et al. (1985), "What is meant is always 'more' than what is said. The cognitive content that directly underlies the speaker's verbal expression is the semantic input" (p. 84). This observation leads to the positing of the pars pro toto ("the part from the whole") principle. People verbalize in pars pro toto fashion, explicitly saying only part of what they "mean." Conversely, comprehension obeys the totum ex parte ("the whole from the part") principle. Acknowledgement of the totum ex parte nature of comprehension is just another way of recognizing the familiar point that people can go beyond the information given (see, e.g., Jacobs, 1985).

Because the PB is always more than the SI, it is plausible to suggest that the PB is transformed into the SI by some sort of selection process. Ideally, this selection will result in a semantic input which is relevant, informative, instrumental, sufficient, truthful, etc. (Grice, 1975). These are abstract, general constraints on the selection process, but "which part of the PB will fulfill these criteria in the case of a concrete utterance depends on the specific features of a situation" (Hoppe-Graff et al., 1985, p. 85).

In summary, the Mannheim group's model of message production posits three steps: (a) the execution of a situation-interpreting program to construct a situation representation embodying both the propositional base of the utterance and specified, situation-specific parameters; (b) pars pro toto selection of semantic input from the propositional base; and (c) low level encoding of semantic input.

## Message Generation and Relevance

It is not possible for a message producer to utter every proposition in his or her currently active propositional base. Consequently, a mechanism is needed to select the subset of propositions that are relevant for expression.

Like Sperber and Wilson (1986), we view the problem of relevance as a matter of explaining what is manifest to speaker and made manifest to hearer; the fundamental principle of message design is "say what's relevant." Whereas Sperber and Wilson (1986) attempt to account for relevance purely in terms of relationships between propositions, we see relevance as a function of focus. And focus originates either in external inputs (which lead to the activation of a unit or pattern of units) or in internal connections between units (when the activation of one unit spreads to an associated unit or pattern of units).

In this view goals are patterns of thoughts about the situation and messages are collations of selected and expressed thoughts. What is most distinctive about this image of message design is the absence of explicit reference to abstract goals or abstract strategies. The question "How do communicators reason from goals to strategies to message contents?" is replaced with "In a given situation, why and how do message producers see different thoughts about the situation as relevant for expression?" And part of the answer is to be found in the connection between patterns of thoughts that represent goals and patterns of message elements.

## A Computer Simulation of Message Generation

We have begun to test this model by writing a computer program to simulate the mapping of goals onto messages (Lambert, 1992). The relationship modelled here is similar to standard plan-based models in that it maps goals onto messages, but there are important differences in how the relationship is implemented. In our model: (a) no explicit logic of action intervenes between goals and messages; (b) message types play no explicit role in processing; (c) goals are represented as distributed patterns of thoughts; and (d) messages are represented as distributed patterns of elemental clauses. Implementing these design features in a PDP framework has enabled us to avoid some of the problems that confront problem-solving models.

PDP networks learn their connection weights by being trained on correct examples of input-output pairs. Thus, to train the model, we developed a method for constructed pairings of activated thoughts and expressed thought units. We asked undergraduate students to respond to a hypothetical broken date situation where an old friend (named Terry) repeatedly cancels dates with the message producer. The message producer is asked to respond to the old friend when s/he calls and asks to reschedule yet another date.

Each subject also completed a Thought Checklist task (Lambert, 1992; O'Keefe, 1992; O'Keefe & Lambert, 1989; Waldron & Cegala, 1992) for this scenario. Responses to the checklist produced a pattern of thoughts represented as a vector of binary values on the 121 variables constituted by the thoughts on the checklist. Each element of each message was classified as representing one of 72 types of thoughts. Based on these codings, messages were represented as vectors of binary values on the 72 variables constituted by the possible content categories.

Lambert (1992) used the example set to train a network using the Quickprop learning algorithm (Fahlman, 1988). Once trained, the network's knowledge was examined by inspecting patterns of exitation and inhibition between thoughts and message elements. Analysis of the network showed that message elements appear to be serving different functions for different message producers. Substantively different (if not completely contradictory) thoughts could strongly influence the same message elements in the same direction (positively or negatively). This finding is puzzling if one believes that messages are planned as unifunctional types that are straightforwardly and statically associated with goals. On such a view, message elements should only be strongly excited by a conceptually coherent set of thoughts. That this is not the case implies that message elements are neither unifunctional nor determinate in meaning. Given a set of training examples drawn from a diverse group of communicators, multiple (possibly contradictory) coalitions of thoughts may excite the same message element because that element is embedded in functionally different webs of meaning for different people.

### Summary

This section described a model of situated message design, where situated message design is conceived as the movement of focus through an organized network of thoughts. The model was designed to avoid the two central problems facing problem-solving models: (a) the problem of instantiating abstract act types as concrete message elements; and (b) the problem of reifying a single functional description of messages. In our model, the instantiation problem never arises, since it employs no abstract act types for computation; instead, thoughts map directly onto concrete message elements. Similarly, it was not necessary to reify a single functional description of messages, since the model needed no explicit logic of action to guide selection of message contents, and since the goal-message mapping is induced from empirical examples.

## A Model of Adaptation and Effects

The previous section provided a view of message structures as expressed thoughts. But a theory of message design connects message structures to message antecedents and message effects. We think of this second component as the theory of adaptation that accompanies the theory of message structure. In this section we develop a model of adaptation in which adaptation is equated with the control of focus by the functional requirements of the current activity, specifically, by an internalized model of the effects of utterance. Our model of adaptation and effects is a developmental one. It reflects the application and extension of ideas first articulated by Rumelhart and Jordan (Jordan, 1989; Jordan & Rumelhart, 1990; Rumelhart, Smolensky, McClelland, & Hinton, 1986).

We first discuss a view of focus as a product of socialization within activities. We then offer a specific model of the processes through which models of activity are internalized and used to guide adaptation. Finally, we discuss the developmental dimension of the model.

### Activity, Relevance, and Focus

Lambert's (1992) model was designed to simulate only one aspect of the process of focusing, the activation of one set of units (message elements) by another set of units (goals, as distributed in patterns of thoughts). But focus depends not only on patterns of connection between units but also on the way input guides focus.

From the standpoint of a theory of communication, one of the most critical ways in which input shapes focus is through transition relevance, i.e., through the way antecedent and projected contributions shape what is relevant to say at any particular juncture. Most current treatments of this issue reflect an essentially Gricean view, in which some form of perspective-taking is seen as driving the calculation of what will be a cooperative contribution (Grice, 1975). For example, on Grice's view, conversationalists are thought to abide by a "quantity" maxim, in which they try to say just what is required and no more in order to meet the listener's needs for information.

Many theorists have attempted to explain cooperation in terms of recipient design, the selection and adaptation of what is said to the requirements of the listener (see, e.g., Clark & Marshall, 1981). However, we see focus as guided not by a model of what the listener requires, but rather by a model of the activity. A study of instruction-giving by Burke (1986) provides a useful illustration of the difference.

Burke (1986) studied the design of instructions given in four media (face-to-face, telephone, audiotape (asynchronous oral), and writing. One group of students, the "experts," was trained to assemble a toy water pump. A second group of students, the "apprentices," received instructions from the experts in how to assemble the pump. Each expert-apprentice pair gave and received instructions in one of the four media. The task thus resembled the one studied by Grosz (1981).

Burke analyzed both the instructions and the behavior of the apprentice in following the instructions. Using an early version of the method we have employed in our recent studies of message organization, Burke segmented the instructions into units that reflected the key steps to be performed in assembling the pump. She then classified these units as to the degree of elaboration of two key themes (specification of the parts to be assembled, specification of the action to be performed in assembly). She found that across media, experts at first organized their instructions in terms of a characteristic pattern of elaboration in which both the designation of the parts and the action description were quite elaborate and detailed. This finding is consistent with research on descriptions summarized earlier; Burke's experts employed a systematic method of navigation through their information about the pump.

However, analysis of the messages in the face-to-face condition showed that in fact, the messages given by the experts were overly informative. In a variety of ways, apprentices signalled that the expert model was not well adapted (for example, they would commonly finish a step before the expert had completed the instructions for it). Over the course of the interaction, experts revised their models and ultimately provided much less detailed directions. In the other three media, where the expert had little access to feedback from the apprentice, no such adaptation of the activity model was observed.

What Burke's findings suggest is that rather than relying on a model of the hearer to compute the informativeness, relevance, or clarity of a contribution, instead speakers have induced models of particular activities. These models are based on direct feedback from hearers, not inferences about what a hearer does or does not know. Because of this, there is no guarantee that a given activity model will in fact embody a correct assessment of the hearer's needs for information. As a consequence, the activity model an individual employs will only be as good as the speaker's range of experience, for only through experience can the speaker receive feedback about the success of a given route to his or her goals.

In summary, then, we see focus as shaped by interactional placement as well as patterns of connection within the field of thoughts that represents a situation. But whereas transition relevance is generally understood as being structured either by sequencing rules that connect types of acts or by recipient design, we see transition relevance as deriving from models of activity that are induced from experience with particular situations. The next section provides a more detailed discussion of activity, adaptation, and message effects.

## Focus, Adaptation, and Effects

The simulation developed by Lambert (1992) represented an attempt to map from current and desired state (represented as a single pattern of thoughts) to a message. Theoretically, a message is chosen because its utterance is expected to transform the current state into the goal state. In symbolic terms:

$$S(I_c, I_g) \rightarrow U_{c,g} \qquad (1)$$

The selection function S() maps the current state $I_c$ and the goal state $I_g$ onto some contextually appropriate utterance $U_{c,g}$. But there is a hidden step here: Why does $S(I_c, I_g)$ map to $U_{c,g}$ and not to some other utterance? The answer is obviously that $U_{c,g}$, when uttered, is expected to transform $I_c$ into $I_g$. Again symbolically, the hidden step involves passing candidate utterances through a model of the world (E), which we call a model of activity:

$$E_c(I_c, U_{c,g}) = I_g \qquad (2)$$

That is, E() takes a message and the current state as input and outputs a description of its effect. With (1) and (2) in mind, a fairly specific account of message generation and its development can be given.

Equation (1) models the message generator; equation (2) models the activity. A complete model of message adaptation and development includes both (see Figure 2). In theory they are connected, with the output of the message generator being the input to the activity model.

$$S(I_c, I_g) \rightarrow U_{c,g} \rightarrow E_c(I_c, U_{c,g}) = I_g \qquad (3)$$

If the current state $I_c$ is factored out, the message generator maps a goal onto an utterance, and the activity model maps a message onto its expected effects:

$$S(I_g) \rightarrow U_g \rightarrow E(U_g) = I_g \qquad (4)$$

It might be noted that this model of message adaptation preserves some of the most important features of problem-solving models (discussed earlier). In particular, the current model accords with the intuition that goals play a key role in message planning and adaptation. But whereas with a problem-solving model, goals are seen as discrete and abstract, within this approach, goals are viewed as situated and distributed in patterns of thoughts. Moreover, a problem-solving approach gives a central role to anticipated effects in the generation process, whereas we conceptualize effects as being primarily involved in processes of reflection and evaluation.

## Learning and Growth

Finally, it is important to note that within this model of adaptation, two key developmental phenomena can be neatly conceptualized: short-term practice effects and long-term acquisition of competence.

Practice effects. It is well known that as individuals become more experienced in an activity, their behavior becomes more practiced and automatic. We recognize the need to account for such practice effects. In terms of the present model, these effects are accounted for in terms of the feedback provided to the message generator by the activity model (Jordan & Rumelhart, 1990).

Message design begins with a pattern of activated thoughts that corresponds to a goal. Activation of the goal state initiates message generation, the selection and transformation of situated knowledge; the goal maps to a message by way of the message generator. Output of the message generator is evaluated by the activity model to determine whether its utterance would bring about the desired effect; if so, the message is uttered. If not, the activity model propagates an error message back through the message generator, changing the weights in that network. The process is repeated until the system is satisfied with the chosen message (perhaps because no more time is available).

Over time, the weights in the message generator will be trained to meet the expectations of the activity model. Hence, reflection and practice should lead to quicker message planning, since fewer iterations should be required to satisfy the activity model.

Acquisition of competence. It is also well established that as individuals mature they develop generally improved skills at message design and adaptation (for a review, see O'Keefe & Delia, 1988). Hence, we also see the need to explain the acquisition of competence over the long term. In the present model, these processes are accounted for in terms of the feedback provided to the activity model by direct experience with the activity. As individuals mature, they accumulate observations of actual effects of their messages; these message-effect pairings are the examples from which the activity model generalizes its knowledge of relations between messages effects in context.

As O'Keefe (O'Keefe, 1988; O'Keefe & Delia, 1989) has argued, many classic progressions in the development of listener-adapted communication can be accounted for in this way, in terms of the acquisition of an increasingly accurate model of communicative activity. As the activity model becomes more accurate, it embodies a better understanding of effects and how to achieve them. To do so requires that the activity model become sensitive to those features of the situation that are objectively, rather than subjectively relevant. These changes, of course, give rise to observed differences in the logic of message design (O'Keefe, 1988).

## Conclusion

The study of message structures has been characterized by a de facto commitment to holism, i.e., by an approach in which messages are treated as instances of acts or strategies. While holism is essential to descriptive models of message meaning, it offers an inadequate approach to causal explanation of message production and effects. A superior causal account can be found in local management approaches to message design (for a more extended discussion of this issue, see Clark, 1989; Lambert, 1992; O'Keefe, 1992).

The approach we have suggested has the advantage of simultaneously acknowledging the utility of holistic-functional analysis and recognizing that any given description of a message is only one of indefinitely many possible true descriptions (O'Keefe, 1987). Any message can be given multiple functional characterizations, some of them mutually incompatible, and all of them equally justified. As can be seen in Brown and Levinson's (1987) analysis of politeness, such a functional description can provide an important and useful perspective on a discourse system.

However, a model that takes such functional descriptions to play a causal role in message generation will necessarily confront not only the instantiation problem we described earlier but also the task of providing a consistent and exhaustive representation of a text. Given the open-endedness and indeterminacy of meaning, such a quest would seem to be bootless.

The local management approach we have developed has given us a distinctive analysis of message structure, one that avoids functional description of message structure in favor of a content-based description of the ideas expressed in the message. In this review we have attempted to show how such a content-based description of message structure, elaborated within a local management view of message planning, can address the enduring problem of explaining message organization and adaptation. By providing independent analyses of message structures and message functions, we make it possible to ask, rather than beg, questions about message design.

References

Agre, P. E., & Chapman, D. (1987). Pengi: An implementation of a theory of activity. In Proceedings of AAAI-87 (pp. 196-201). Los Altos, CA: Morgan Kaufmann.

Agre, P. E., & Chapman, D. (1990). What are plans for? Robotics and Autonomous Systems, 6, 17-34.

Anderson, J. A. , & Hinton, G. E. (1981). Models of information processing in the brain. In G. E. Hinton and J. A. Anderson (Eds.), Parallel models of associative memory (pp. 9-48). Hillsdale, NJ: Lawrence Erlbaum Associates.

Appelt, D. E. (1982). Planning natural-language utterances. In Proceedings of the National Conference on Artificial Intelligence (pp. 59-62). Los Altos, CA: Morgan Kaufman.

Burke, J. A. (1986). Interacting plans in the accomplishment of a practical activity. In D. G. Ellis & W. A. Donahue (Eds.), Contemporary issues in language and discourse processes (pp. 203-222). Hillsdale, NJ: Erlbaum.

Chafe, W. L. (1979). The flow of thought and the flow of language. In T. Givon (Ed.), Syntax and semantics, Vol. 12: Discourse and syntax (pp. 159-182). New York: Academic.

Clark, A. (1989). Microcognition: Philosophy, cognitive science, and parallel distributed processing. Cambridge, MA: MIT Press.

Clark, H., & Marshall, C. R. (1981). Definite reference and mutual knowledge. In A. K. Joshi, B. L. Webber, & I. A. Sag (Eds.), Elements of discourse understanding (pp. 10-63). Cambridge: Cambridge University Press.

Cohen, P. R., & Perrault, C. R. (1979). Elements of a plan-based theory of speech acts. Cognitive Science, 3, 177-212.

Dell, G., & Juliano, C. (1990). Connectionist approaches to the production of words. Unpublished manuscript, Department of Psychology, University of Illinois, Urbana-Champaign.

Dennett, D. (1987). The intentional stance. Cambridge, MA: MIT Press.

Fahlman, S. E. (1988). An empirical study of learning speed in back-propagation networks (Tech. Rept. No. CMU-CS-88-162). Pittsburgh, PA: Carnegie Mellon University.

Fikes, R. E., & Nilsson, N. J. (1971). STRIPS: A new approach to the application of theorum proving to problem solving. Artificial Intelligence, 2, 189-208.

Garfinkel, H. (1967). Studies in ethnomethodology. Cambridge: Polity Press.

Garrett, M. F. (1975). The analysis of sentence production. In G. Bower (Ed.), Psychology of learning and motivation (pp. 133-177). New York: Academic Press.

Garrett, M. F. (1988). Processes in language production. In F. J. Newmeyer (Ed.), Linguistics: The Cambridge survey, vol. 3: Language: Psychological and sociological aspects (pp. 69-96). Cambridge: Cambridge University Press.

Green, G. (1989). Pragmatics and natural language understanding. Hillsdale, NJ: Erlbaum.
Greene, J. O. (1989). Action assembly theory: Metatheoretical commitments, theoretical propositions, and empirical applications. In B. Dervin, L. Grossberg, B. J. O'Keefe, & E. Wartella (Eds.), Rethinking Communication, Vol. 2: Paradigm exemplars (pp. 117-128). Newbury Park, CA: Sage.
Grice, H. P. (1975). Logic and conversation. In P. Cole & J. L. Morgan (Eds.), Syntax and semantics, vol. 3: Speech acts. New York: Academic.
Grosz, B. J. (1981). Focusing and description in natural language dialogues. In A. Joshi, B. Webber, and I. Sag (Eds.), Elements of discourse understanding (pp. 84-105). Cambridge: Cambridge University Press.
Hermann, T. (1983). Speech and situation: A psychological conception of situated speaking. Berlin: Springer-Verlag.
Hillocks, G. (1986). Research on written composition: New directions for teaching. Urbana, IL: NCTE.
Hinton, G. (1990). Connectionist learning procedures. In J. Carbonell (Ed.), Machine learning: Paradigms and methods (pp. 185-234). Cambridge, MA: MIT Press.
Hoppe-Graff, S., Herrmann, T., Winterhoff-Spurk, P., & Mangold, R. (1985). Speech and situation: A general model for the process of speech production. In J. Forgas (Ed.), Language and social situations (pp. 81-95). New York: Springer-Verlag.
Hovy, E. (1990). Pragmatics and natural language generation. Artificial Intelligence, 43, 153-197.
Jacobs, S. (1985). Language. In M. L. Knapp & G. R. Miller (Eds.), Handbook of interpersonal communication. Newbury Park, CA: Sage.
Jordan, M. (1989). Action. In M. I. Posner (Ed.), Foundations of cognitive science (pp. 727-767). Cambridge, MA: MIT Press.
Jordan, M., & Rumelhart, D. E. (1990). Forward models: Supervised learning with a distal teacher. Occasional paper No. 40, Center for Cognitive Science, MIT, Cambridge, MA.
Kellermann, K., & Lim, T.-S. (1989). Conversational acquaintance: The flexibility of routinized behaviors. In B. Dervin, L. Grossberg, B. J. O'Keefe, & E. Wartella (Eds.), Rethinking Communication, Vol. 2: Paradigm exemplars (pp. 172-187). Newbury Park, CA: Sage.
Kirsch, D. (1988). Competence models are causal. Behavioral and Brain Sciences, 11, 515-517.
Lambert, B. L. (1992). A connectionist model of message design. Unpublished doctoral dissertation, University of Illinois, Urbana-Champaign.
Lambert, B. L., & Lee, J. Y. (1993). Patient perceptions of pharmacy students' hypertension compliance gaining messages: Effects of message design logic and content themes. Unpublished paper, Department of Pharmacy Administration, University of Illinois, Chicago.
Lambert, B. L., & O'Keefe, B. J. (1990, November). Modelling message design. Paper presented at the annual meeting of the Speech Communication Association, Chicago, IL.
Lave, J. (1988). Cognition in practice: Mind, mathematics and culture in everyday life. Cambridge: Cambridge University Press.
Levelt, W. J. M. (1989). Speaking: From intention to articulation. Cambridge, MA: MIT Press.
Levinson, S. C. (1982). Pragmatics. Cambridge: Cambridge University

Press.
Levy, D. M. (1979). Communicative goals and strategies: Between discourse and syntax. In T. Givon (Ed.), Syntax and semantics, Vol. 12: Discourse and syntax (pp. 183-212). New York: Academic.
Linde, C., & Labov, W. (1975). Spatial networks as a site for the study of language and thought. Language, 51, 924-929.
Marwell, G., & Schmitt, D. R. (1967a). Compliance-gaining behavior: A synthesis and model. Sociological Quarterly, 8, 317-328.
Marwell, G., & Schmitt, D. R. (1967b). Dimensions of compliance-gaining behavior: An empirical analysis. Sociometry, 30, 3520-364.
McLaughlin, M. (1984). Conversation. Newbury Park, CA: Sage.
McClelland, J. L., & Rumelhart, D. E., & the PDP Research Group. (1986). Parallel distributed processing: Explorations in the microstructure of cognition, Vol. 2: Psychological and biological models. Cambridge, MA: MIT Press.
McClelland, J. A., Rumelhart, D, & Hinton, G. E. (1986). The appeal of PDP. In D. Rumelhart, J. L. McClelland, & the PDP Research Group (Eds.), Parallel distributed processing: Explorations in the microstructure of cognition (Vol. 1, pp. 3-44). Cambridge, MA: MIT Press.
Mikkulainen, R. (1993). Subsymbolic natural language processing: An integrated model of scripts, lexicon, and memory. Cambridge, MA: MIT Press.
Miller, G., Galanter, E., & Pribram, K. (1960). Plans and the structure of behavior. New York: Holt, Rinehart and Winston.
Minsky, M., & Papert, S. (1969). Perceptrons. Cambridge, MA: MIT Press.
Newell, A. (1980). Physical symbol systems. Cognitive Science, 4, 135-183.
Newell, A., & Simon, H. (1972). Human problem solving. Englewood Cliffs, NJ: Prentice-Hall.
O'Keefe, B. J. (1988). The logic of message design. Communication Monographs, 55, 80-103.
O'Keefe, B. J. (1991, May). Managing constraints on message design: The relation between social structures and communicative forms. Paper presented at the annual meeting of the International Communication Association, Chicago.
O'Keefe, B. J. (1992). Developing and testing rational models of message design. Human Communication Research, 18, 637-649.
O'Keefe, B. J. (in press). Variation, adaptation, and functional explanation in the study of message design. In G. Philipsen & T. Albrecht (Eds.), Developing theories of communication. Albany, NY: SUNY Press.
O'Keefe, B. J., & Delia, J. G. (1988). Communicative tasks and communicative practices: The development of audience-centered message production. In B. A. Rafoth & D. L. Rubin (Eds.), The social construction of written communication (pp. 70-98). Norwood, NJ: Ablex.
O'Keefe, B. J., & Lambert, B. L. (1989, November). Effects of message design logic on the communication of intention. Paper presented at the annual meeting of the Speech Communication Association, San Francisco.
O'Keefe, D. J. (1987). Describing messages. Paper presented at the annual meeting of the Speech Communication Association, Boston, MA.

Roger, D., & Bull, P. (1987). Conversation: An interdisciplinary perspective. Clevedon: Multilingual Matters.
Rumelhart, D.E., Smolensky, P., McClelland, J. L., & Hinton, G. (1986). Schemata and sequential thought processes in PDP models. In J. L. McClelland & D. E. Rumelhart (Eds.), Parallel distributed processing: Explorations in the microstructure of cognition, Vol. 2: Psychological and biological models (pp. 7-57). Cambridge, MA: MIT Press.
Rumelhart, D. E., & McClelland, J. (1986). Parallel distributed processing: Explorations in the microstructure of cognition, Vol. 1: Foundations. Cambridge, MA: MIT Press.
Sacerdoti, E. (1974). Planning in a hierarchy of abstraction spaces. Artificial Intelligence, 5, 115-135.
Saeki, M., & O'Keefe, B. J. (1993). Refusals and rejections: Designing messages to serve multiple goals. Paper presented at the annual meeting of the Speech Communication Association.
Schegloff, E. (1968). Sequencing in conversational openings. American Anthropologist, 70.
Searle, J. R. (1969). Speech acts. Cambridge: Cambridge University Press.
Seibold, D. R., Cantrill, J. G., & Meyers, R. A. (1985). Communication and interpersonal influence. In M. L. Knapp & G. R. Miller (Eds.), Handbook of interpersonal communication (pp. 551-611). Beverly Hills, CA: Sage.
Sibun, P. (1990). The local organization of text. In Proceedings of the Fifth International Workshop on Natural Language Generation (pp. 120-127). Los Altos, CA: Morgan Kaufman.
Sibun, P. (1991a). Salix: A strategy-based architecture for text organization. Unpublished manuscript, Xerox Palo Alto Research Center, Palo Alto, CA.
Sibun, P. (1991b). Locally organized text generation. Unpublished doctoral dissertation, University of Massachusetts, Amherst.
Sibun, P. (in press). Generating text without trees. Computational Intelligence, 8.
Smolensky, P. (1988). On the proper treatment of connectionism. Behavioral and Brain Sciences, 11, 1-59.
Sperber, D., & Wilson, D. (1986). Relevance: Communication and cognition. Cambridge, MA: Harvard University Press.
Suchman, L. (1987). Plans and situated actions: The problem of man-machine communication. Cambridge: Cambridge University Press.
Tannen, D. (1993). The relativity of linguistic strategies; Rethinking power and solidarity in gender and dominance. In D. Tannen (Ed.), Gender and conversational interaction (pp. 165-188). Oxford: Oxford University Press.
Taylor, T. J., & Cameron, D. (1989). Analysing conversation: Rules and units in the structure of talk. Oxford: Pergamon Press.
Ullmer-Ehrich, V. (1982). The structure of living space descriptions. In R. J. Jarvella & W. Klein (Eds.), Speech, place, and action: Studies in deixis and related topics (pp. 219-249). New York: Wiley.
Waldron, V., & Cegala, D. (1992). Assessing conversational cognition: Levels of cognitive theory and associated methodological requirements. Human Communication Research, 18, 599-622.
Wunderlich, D., & Reinelt, R. (1982). How to get there from here. In R. J. Jarvella & W. Klein (Eds.), Speech, place, and action: Studies in deixis and related topics (pp. 183-201). New York: Wiley.

Figure 1.  Herrmann's Speech Production Model

```
ZDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDD?
3                                                                      3
3                  ZDDDDDDDDDDDDDDDDDDDDDDDDDDDDD?                     3
3                  3                             3                     3
3   DATA BASE:     3        SITUATION            3                     3
3                  @DDDDDDDDDDDDDDDBDDDDDDDDDDDDDY                     3
CDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDEDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDD4
3                  ZDDDDDDDDDDDDDDDADDDDDDDDDDDDD?   ZDDDDDDDDDDDDDDD?3
3   PROGRAM:       3  Situation-interpreting     CDDD4 LTM:          33
3                  3  program(s): Construction   3   3 Procedural &  33
3                  3  of the cognitive repre-    CDDD4 declarative   33
3                  3  sentation of the situation 3   3 knowledge     33
3                  @DDDDDDDDDDDDDDDBDDDDDDDDDDDDDY   @DDDDDDDDDDDDDDDY3
3                                  3                                   3
3                  ZDDDDDDDDDDDDDDDADDDDDDDDDDDDD?                     3
3   DATA BASE:     3  Situation representation:  3                     3
3                  3      ~             ~        3                     3
3                  3 ZDDDDADDDDD?ZDDDDDDADDDDDDD?3                     3
3                  3 3Specified 33 Propositional33                     3
3                  3 3situation 33 base PB:     33                     3
3                  3 3parameters33 Active verbal33                     3
3                  3 3          33 knowledge    33                     3
3                  3 @DDDDDDDDDDY@DDDDDDDDDDDDDDY3                     3
3                  @DDDDDDBDDDDDDDDDDDDDDBDDDDDDDY                     3
3                         3              3                             3
3                  ZDDDDDDADDDDDDDDDDDDDDADDDDDDD?                     3
3                  3                             3                     3
3   PROGRAM:       3  Selection program:         3                     3
3                  3  Selection of subset of PB  3                     3
3                  @DDDDDDDDDDDDDDDBDDDDDDDDDDDDDY                     3
3                  ZDDDDDDDDDDDDDDDADDDDDDDDDDDDD?                     3
3   DATA BASE:     3  Semantic input  SI:  Subset3                     3
3                  3  of PB                      3                     3
3                  @DDDDDDDDDDDDDDDBDDDDDDDDDDDDDY                     3
3                  ZDDDDDDDDDDDDDDDADDDDDDDDDDDDD?                     3
3   PROGRAM:       3  Encoding program(s):       3                     3
3                  3  Syntactic, lexical, and    3                     3
3                  3  prosodic encoding of SI    3                     3
3                  @DDDDDDDDDDDDDDDBDDDDDDDDDDDDDY                     3
3                                  3                                   3
CDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDADDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDD4
3                                                                      3
3                     Observable utterance                             3
3                                                                      3
@DDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDY
```

Figure 2.  Relationships among message generator, activity model, and activity.  Framework patterned after Jordan and Rumelhart (1990).

```
                   ZDDDDDDDDDDDDDDD?    ZDDDDDDDDDDDDDDD?
                   3 Current state 3    3 Goal state    3
                   @DDBDDDDBDDDDDDDY    @DDDDDDDBDDDDDDDY
                      3    3                    3
```

```
                        3  ZDADDDDDDDDDDDDDDDDDDDDDDDAD?
        ZDDDDDDDDDDDDDDDDEDD4  Message Generator     3
        3               3  @DDDDDDDDDDDDDDDDDDDDDDDDDDY
        3               3
        3               3
ZDDDDDDDADDDDDDD?    ZDDADDDDDDDDDDDD?    ZDDDDDDDDDDDDDDD?  ZDDDDDDDDDDD?
3Feedback from  3    3 Current state 3    3  Message     CDD4Utterance 3
3Activity Model 3    @DDDDDDDDDDBDDDDY    @DDDDBDDDDDDDDDY  @DDDDDDDDDDDY
@DDDDDDDBDDDDDDDY              3              3
        3                  ZDDDDADDDDDDDDDDDDDDADDDD?   ZDDDDDDDDDDDDDDD?
        @DDDDDDDDDDDDDDDDDDDD4                        CDDD4Feedback from 3
                           3    Activity Model      3   3Activity      3
                           @DDDDDDDDDDDBDDDDDDDDDDDDDY   @DDDDDDDDDDDDDDDY
                                       3
                           ZDDDDDDDDDDDDADDDDDDDDDDDDD?
                           3  Predicted Next State  3
                           @DDDDDDDDDDDDDDDDDDDDDDDDDDY
```