© Copyright by Bruce Laurence Lambert, 1992

A CONNECTIONIST MODEL OF MESSAGE DESIGN

BY

BRUCE LAURENCE LAMBERT

A.B., University of Illinois at Urbana-Champaign, 1987
A.M., University of Illinois at Urbana-Champaign, 1988

THESIS

Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Speech Communication
in the Graduate College of the
University of Illinois at Urbana-Champaign, 1992

Urbana, Illinois

A CONNECTIONIST MODEL OF MESSAGE DESIGN

Bruce Laurence Lambert, Ph.D.
Department of Speech Communication
University of Illinois at Urbana-Champaign, 1992
Barbara J. O'Keefe, Advisor

Complexity in plan-based models of language production has typically been managed by the systematic exploitation of abstractions. Most significantly, abstract goals replace concrete situations, and act types replace instantiated linguistic contents. Consequently, plan-based models involve three key transformations: (a) the mapping of concrete situations onto abstract goals, (b) the mapping of abstract goals onto act types, and (c) the mapping of act types onto linguistic tokens. Transformations are made possible by the prior functional indexing of situations, acts and utterances. Situations are indexed by the goals they embody. Acts are indexed by the goals they achieve, and linguistic types are indexed by the acts they instantiate. However, such functional indexing schemes are based on untenable assumptions about language and communication. Communicative means-ends relations are not fixed, and the assignment of functional significance to decontextualized linguistic forms contradicts a substantial body of evidence about the contextual dependency of language function.

A connectionist model of message production was proposed that avoids these difficulties. The model used Quickprop, a modified version of the back-propagation learning algorithm, to learn a mapping from thoughts to the elements of a phrasal lexicon for a hypothetical broken date situation. Goals were represented as distributed patterns of context-specific thoughts. Messages were represented as patterns of phrases. Three studies were conducted to develop and test the model. The first study used protocol analysis to validate a checklist method for thought elicitation. The second study described a prototype system for automatically coding independent clauses into phrase categories (called message frames), and the system's ability to correctly classify new examples was found to be quite limited. The final study analyzed the performance of the connectionist model. The generalization performance of the model was limited by the small number of examples, but an analysis of the receptive fields of several message frames was consistent with previous research about individual differences in reasoning about communication and supported the assertion that individual linguistic forms are multifunctional. The model provided preliminary support for a situated, phrasal-lexical model of message design that does not rely on decontextualized, functionally indexed abstractions.

ACKNOWLEDGMENTS

I would not have been able to complete this project were it not for the support, kindness, and generosity of many teachers, colleagues, and friends. I would like to acknowledge the contribution of the members of the Neural Net Group of the Center for Complex Systems Research at the University of Illinois, especially Subutai Ahmad, Bruce Katz, Pankaj Mehra, Bartlett Mel, Klaus Obermayer, and Steve Omohundro, all of whom put up with my neural naivete and encouraged me to pursue ideas that were at the time extremely crude.

I am grateful also to Scott Fahlman and Terry Regier for placing their versions of the Quickprop algorithm into the public domain and for answering a variety of technical questions. I have benefitted also from discussions with Scott Jacobs and Penni Sibun about the nature of language production, and Bob Mrtek provided useful insight into a couple of statistical analysis issues.

Marcy Dorfman provided key guidance and direction in the earliest stages of this project, in spite of her substantive disagreements with my claims. In addition, both Gary Dell and Chip Bruce listened patiently and provided extremely helpful feedback and perspective on my arguments and methods as they developed. To Stephen Lu and all of the members of the Knowledge-Based Engineering Systems Research Lab, I am grateful. The lab members provided hours of technical support for a UNIX novice, and Stephen Lu generously made the extensive computing resources of his lab available so that I could complete the modelling part of this project.

I am particularly indebted to Bill Brewer, who taught me about knowledge, and to Andy Pickering, who taught me about practice. Both provided excellent companionship as well as advice and instruction, and both made graduate school a far richer experience than it would have been otherwise. This thesis is in part an attempt to reconcile their perspectives, and for that reason it will likely leave both of them less than completely satisfied.

I want to thank Dan O'Keefe and Noshir Contractor for serving on my committee, and for working with me before and during my thesis experience. Dan O'Keefe especially exemplified for me what it meant to think carefully about communication theory, and my hope is that he finds my thinking not overly reckless.

For helping me stay sane during the composition of this document, I thank Eric Lambert, my brother, and Kevin Wolff, who is like a brother to me. For the same reason, I thank Mike Pruetting and Al Poeppel. Many a thesis worry was forgotten in their company.

I want to thank David Lambert, my other brother, for introducing me to the idea of parallel distributed processing in 1987, before I even entered graduate school. More than that, he has taught me most of what I know about neural networks and much of what I know about computers and computer programming. He has been and continues to be a source of intellectual guidance and support without peer. Tremendous thanks are also owed to David Tcheng, a long-time friend and collaborator whose knowlege of artificial intelligence, LISP, and machine learning has been an invaluable resource to me, a resource that I have drawn on too many times to mention.

Above all, Barbara O'Keefe has my appreciation and admiration. This thesis is the end product of five years of intense discussion and debate with her about language, action, meaning, context, etc. It is no exaggeration to say that the ideas that populate this thesis literally never would have occurred to me if not for her constant attempts to help me understand. Her influence on me is immeasurable. Even my most sublime moments of inspiration often appeared, on further consideration, to be instances when something Barb

had said months or years ealier finally made sense to me. I am fortunate to count her among my friends.

Finally, I thank my parents and my wife Naomi. Were it not for their boundless support, I surely would not have had the strength or desire to finish this project, and perhaps not even to begin it.

TABLE OF CONTENTS

Page

CHAPTER ONE: INTRODUCTION 1

Two Models of Intelligent Adaptation
Rationality as Generative Mechanism: The Planning Model
Rationality as Situated Achievement: Local Management..
Summary
Primary Research Questions.
Outline of Chapters

CHAPTER TWO: COMPREHENSIVE MODELS OF LANGUAGE PRODUCTION 7

Overview
Three Processes in Language Production
Conceptualization
Formulation
Articulation
Representation and Processing in Message-Level Production
Knowledge Representation
Plan-Based Models of Message Production
AI, Problem Solving, and Planning
Speech Acts and Plans
Speech Act Models of Message Planning and Production
Functional Abstraction and Plan-Based Models.
Functional Indexing Schemes
Summary

CHAPTER THREE: A CRITIQUE OF TYPIFICATION IN COMMUNICATION THEORY 18

The Typification of Situations
The Typification of Messages
Standard Models Depend on the Premises
Validity of the Premises.
The Indefinite Describability of Messages
The Typification of Means-Ends Relations.
Functional Message Production
Examples of Functional Diversity in Message Designs.
Typification and Message Diversity
Summary

CHAPTER FOUR: LOCAL MANAGEMENT AND MESSAGE DESIGN: A PDP VIEW 32

A Situation-Based Representation of Message Structure
Distributing the Representation of Communicator Goals
A Model of Message Production
Herrmann's Basic Model of Speech Production
Selection of Semantic Input
Summary

CHAPTER FIVE: DEVELOPMENT AND VALIDATION OF THE THOUGHT CHECKLIST 37

Methods and Design
Subjects
Thought Checklist Condition
Protocol Condition
Analysis and Results
Quantity of Reported Thoughts
Representativeness of the Thought Checklist
Comprehensiveness of the Thought Checklist
General Discussion
Summary

CHAPTER SIX: DEVELOPMENT AND VALIDATION OF THE MESSAGE FRAME CODING SYSTEM. 44

Overview
Representation of the Broken Date Messages as an Activation Vector
Redundancy Within Message Frames
Reliability
AutoCoder: Development of a Prototype Automated Coding System
Messages
Message Elements
Keywords
Phrases
Frames
The Auto-Coder Algorithm
Words-to-Keywords
Keywords-to-Phrases and Phrases-to-Frames
Quantitative Characteristics of Keyword and Phrase Definitions
Performance on Unseen Cases.

Summary

CHAPTER SEVEN: A CONNECTIONIST MODEL OF CONTENT SELECTION IN MESSAGE DESIGN 60

Overview
Related Models of Language Processing
PDP Models
Local Management Models
Review of Theoretical Model
Methods
Subjects
Representation of Input Features
Representation of Output Features
The Learning Algorithm
Speeding Up Back-Propagation:
The Quickprop Algorithm
Architecture of the Model
Network Performance
Analysis of Weights
Methods
Analysis of Frame 54
Analysis of Frame 49
Analysis of Frame 65
Analysis of Frame 6
Analysis of Frame 1
Goals as Distributed Patterns of Thoughts
Clustering Method
Potential for the Study of Individual Differences
Summary

CHAPTER EIGHT: CONCLUSION 89

Overview
Summary and Interpretation of Findings.
Functional Abstraction and Plan-Based Models
Message Production Without Types.
A Checklist Method for Thought Elicitation.
Semi-Automated Message Coding
Mapping Thoughts onto Messages
Limitations of the Current Investigation
Limitations of the Thought Checklist Task
Limitations of the Message Frame Coding System
Limitations of the PDP Model

Future Work
Implications and General Discussion
Content-Based Models of Message Production

REFERENCES 99

APPENDIX A 107

APPENDIX B 115

APPENDIX C 117

APPENDIX D 120

VITA 170

LIST OF TABLES

Table		Page
1	Recommended Additions to the Thought Checklist	42
2	Frequency and Classification Data for 72 Message Frames	45
3	Selected Message Frames Illustrating Within-Frame Redundancy	48
4	Message Elements at Three Levels of Abstraction	54
5	Examples of Keyword-to-Phrase Misclassification	57
6	Performance Measures for Best Network Architectures	67
7	Number of Epochs to Learn Training Set Perfectly	68
8	Comparison of Predicted Versus Actual Message Frames and Model-Generated Versus Actual Message Instances	70
9	Excitatory and Inhibitory Receptive Fields for Frame 54	74
10	Excitatory and Inhibitory Receptive Fields for Frame 49	76
11	Excitatory and Inhibitory Receptive Fields for Frame 65	77
12	Excitatory and Inhibitory Receptive Fields for Frame 6	79
13	Excitatory and Inhibitory Receptive Fields for Frame 1	80
14	Thought Clusters with Conceptual Labels	82
15	Cluster Summary for 28 Thought Clusters	87

LIST OF FIGURES

Figure		Page
1	A comprehensive model of language production	
2	Four mappings of interest to communication theorists	
3	Herrmann's basic model of speech production	
4	Pseudo-code for WORD-TO-KEYWORD classification	
5	Schematic model of thought-message mapping	
6	Frequency distribution of message frames	
7	Relationship between raw accuracy scores and frame frequency for three measures of accuracy	
8	Relationship between transformed accuracy scores and frame frequency for three measures of accuracy	
9	Relationship between message generator, world model, and world	

CHAPTER ONE
INTRODUCTION

> We identify strategic message construction as the key locus of the interface of language and society. (Brown & Levinson, 1987, p. 56)

Several general frameworks have been offered for understanding the process of message production (Butterworth, 1980; van Dijk & Kintsch, 1983; Garrett, 1975, 1988; Greene, 1984a; Herrmann, 1983; Kellerman, Broetzman, Lim, & Kitao, 1989; Levelt, 1989; O'Keefe, 1988; Schlesinger, 1977). Though these approaches differ in their details, they agree in viewing message production as a rational process of reasoning from goals to linguistic means for achieving those goals (Brown & Levinson, 1987; Bruce, 1986; Clark & Clark, 1977; Levelt, 1989; O'Keefe, 1988, 1990a, 1990b). Since message production is understood as a cognitive process that maps goals to messages, it is important to characterize that process in detail. Hence, the central task in the study of message production is to explain how people get from goals to messages.

By far the most popular and fruitful approach to this task has been to suggest that goals are transformed into messages by a process of rational planning (Appelt, 1986; Bruce, 1980, 1986; Clark & Clark, 1977; Cohen, Morgan, & Pollack, 1990; Cohen & Perrault, 1979; Hovy, 1990; Miller, Galanter, & Pribram, 1960; Schank & Abelson, 1977). The planning model has been extremely successful in extending our understanding of the process of message production. The work generated from this perspective continues to be the most detailed and formal work available on the subject. However, there is an increasingly convincing case being made that the basic assumptions, methods, and mode of explanation used in plan-based models of message production need to be reconsidered (Agre, 1990, in press; Agre & Chapman, 1987, 1990; A. Clark, 1989; Goffman, 1981; Lave, 1988; Levinson, 1979a; O'Keefe, 1990a, 1990b; Sibun, 1991b; Suchman, 1987). These concerns justify both a careful reexamination of plan-based approaches to message production and a preliminary search for alternatives. In the remainder of this introductory chapter, I compare and contrast two theories of intelligent adaptation, one based on planning and one based on locally managed, situated action. The chapter ends with a statement of research questions and an overview of the thesis.

Two Models of Intelligent Adaptation

Many problems in cognitive science are problems of adaptation (Anderson, 1983; Rumelhart & McClelland, 1986). Those who study the production of behavior are especially interested in how action is flexibly adapted to changing tasks and circumstances. The basic outlook adopted here views message production as a problem of adaptation.

This is hardly a novel suggestion. One persistent claim about message production has been that messages are successful just to the extent that they are well adapted to the wants, needs, knowledge, or general perspective of the addressee (Flavell, 1968; O'Keefe & Delia, 1982, 1985; Selman, 1980; Shantz, 1975). Thus, each theory of message production relies more or less self-consciously on a model of intelligent adaptation. Contemporary social science provides at least two alternative models of adaptation to undergird theories of message production: one based on plans (Newell & Simon, 1972), and one based on locally managed, situated actions (Garfinkel, 1967; Lave, 1988; Suchman, 1987). Each view is summarized below.

Rationality as Generative Mechanism: The Planning Model

Planning and standard cognitivism. The first model I will call the planning model. The

planning model treats adaptation as a rational process of reasoning from ends to means. Plans are typically stated in a vocabulary of decontextualized act-types, and they are constructed by the deployment of abstract strategies. The relation between plan and action is seen as causal and deterministic. Plans exist prior to actions and determine their course (Suchman, 1987). The modern planning model originated out of cognitive psychology (Miller, Galanter, & Pribram, 1960; Newell & Simon, 1972). The vast majority of theories of message production are based on the planning model (Appelt, 1986; Bruce, 1986; Cohen & Perrault, 1979; Levelt, 1989).

A commitment to planning as the best explanation of intelligent adaptation is the hallmark of what we will call standard cognitivism (A. Clark, 1989; Suchman, 1987). Standard cognitivism sees intelligent action as the result of symbolic computation in the language of thought. The language of thought is a symbolic representational scheme with compositional syntax and semantics, and is, in that sense, isomorphic with natural languages (Fodor & Pylyshyn, 1988). The language of thought is said to be implemented in a physical symbol system. The brain is one such system, capable of creating, storing, comparing, and moving symbolic tokens (Newell, 1980).

Problem solving is the primary mechanism supporting intelligent action on this view. For a given problem, computations in the language of thought make it possible to construct a representation of the problem known as a problem space. A problem space is a symbolic representation of all possible states of the problem. The two most important states in the problem space are the current state and the goal state. The task of the problem solver is always to get from the current state to the goal state. Problem states change as a consequence of the application of operators, which are abstract actions. Thus, problem solving consists in a search for a sequence of acts that will transform the current state to the goal state (Newell & Simon, 1972). The close connection between the planning model and standard cognitivism is important. It suggests that they stand or fall together. If the planning model turns out to be an inadequate model of intelligent adaptation, serious doubts are cast on standard cognitivism as a model of mind.

Problems with the planning model. Despite the success of the planning model at modelling some cognitive processes, there are nagging problems. Simply put, the planning model ignores the contingency of action in the world. The material and social worlds often resist our attempts to act upon them (Pickering, 1990). The planning model does not respect the embodied nature of action in the world, preferring instead to restrict action within abstract, symbolic problem domains (Lave, 1988). Furthermore, plans underspecify the actions they supposedly determine, and the goals which guide the construction of plans are often not known in advance (Keller & Keller, 1988; Pickering, in press; Suchman, 1987). Plans describe sequences of abstract act-types, but provide little insight into how act-types are instantiated as embodied actions. Recognizing these shortcomings in traditional planning models, one could proceed in two different directions. One might attempt to improve on the planning model (for a commentary on such attempts, see Agre, 1990; Agre & Chapman, 1990; Miller, 1990). Or one might begin to question the validity of the planning model as a basis for understanding intelligent adaptation. The next section describes the work of those who have begun to seek an alternative to the planning model.

Rationality as Situated Achievement: Local Management

Definition of local management. The second model of intelligent adaptation will be referred to as the local management model (a.k.a., the situated action model). The local management model has its roots in sociology (Garfinkel, 1967; Heritage, 1984; Sacks, Schegloff, & Jefferson, 1974), but similar themes have been developed in cognitive anthropology (Lave, 1988; Rogoff & Lave, 1984; Suchman, 1987; Wertsch, 1985), social

studies of science (Pickering, 1990, in press), and artificial intelligence (Agre, 1987; Agre & Chapman, 1990; Chapman, 1987; Sibun, 1990, 1991a, 1991b, in press). The local management model sees adaptation as a continuous process of reacting to the concrete contingencies of sequentially organized, situated action. On this view, abstract act-types and global organizational structures (like plans) may be useful as post-hoc, summary descriptions of adaptive behavior, but adaptation itself is made possible by processes not easily described in the everyday language of plans and goals (A. Clark, 1989; Suchman, 1987). Thus where the planning model is abstract, disembodied, and decontextualized, the local management (or situated action) model is concrete, embodied, and contextual.

Elaborating concepts of local management. The concept of local management has received its fullest expression in conversation analysis. Turn-taking, adjacency pair organization, and conversational repair mechanisms have all been described in terms of local management (Levinson, 1983). In each case, analysts note that the appearance of coherence and stability is an achievement embedded in the sequential give-and-take of interaction. As suggestive as these analyses have been, they have dealt primarily with the local management of access to the floor, and less with the local organization of expressed content.

Content-based models of production. If local management is to be a credible alternative to the planning model of adaptation, it is vital to develop more elaborated accounts of the local organization of expressed content. Sibun's (1991b) work on the local organization of text provides an example of how one might go about extending the concept of local management to include content as well as floor management. Rather than generating texts from formal structures (e.g., plans or discourse schemas), Sibun (1991a) argues that coherent texts can be produced by "exploiting the rich local structure of the subject matter" (p. 2). While Sibun was able to model locally managed descriptions using standard, symbolic artificial intelligence (AI) tools, there is reason to believe that an alternative computational strategy might be useful for modelling more complex communicative tasks involving multiple goals. Before closing this discussion, some possible connections between local management and newer "connectionist" or parallel distributed processing models are examined.

Connectionism and local management. As I stated above, the planning model is inextricably bound up with standard cognitivist assumptions about the nature of intelligent action. I also argued that if the planning model fails, so fails standard cognitivism, at least in some respects. Given the problems now facing the planning model, it is appropriate to look for a new kind of cognitivism to go along with local management conceptions of situated action. As advocates of local management have pressed the case against plan-based models of adaptation, an alternative to standard cognitivism has begun to emerge.

The challenge to standard cognitivism and the physical symbol system hypothesis (Newell, 1980) comes from an interdisciplinary community advocating a parallel distributed processing (henceforth PDP) approach to cognition (A. Clark, 1989; Hanson & Burr, 1990; Levine, 1991; McClelland & Rumelhart, 1986; Rumelhart & McClelland, 1986; Smolensky, 1988). Rather than symbols and search, the basic building blocks in the PDP framework are units and connections.

PDP models are highly interconnected networks of simple computing elements. Each element, or unit, has a numeric activation value. Units influence one another by passing activation along these weighted connections. The effect of some unit *i* on another unit *j* is a function of the product of *i*'s activation value and the weight connecting *i* and *j*. In many PDP models, units are divided into input and output layers. Information is represented in the network as patterns of activation over semantically interpreted pools of units. The activation of a unit typically corresponds to the degree of presence or absence of some semantic feature

of the input or output representation. The output resulting from a particular input is a function of the weights on the intervening connections.

The standard method of cognitive modelling consists in constructing computer programs that mimic the input/output characteristics of some aspect of human cognitive functioning. One attractive feature of PDP models is the existence of procedures for modifying connection weights such that they approximate the input/output profile exemplified in a set of "correct" input/output examples. Several such weight modification procedures exist (see Hanson & Burr, 1990; Hinton, 1990). The next section argues that PDP models are better suited to modelling local management processes than are standard, symbolic approaches (A. Clark, 1989; Smolensky, 1988).

Descriptive and causal cognitive science. To understand why PDP models are better equipped to handle local management processes, it is helpful to consider the distinction between intentional stance, design stance, and physical stance explanations of behavior (Dennett, 1987) and the related distinction between descriptive and causal cognitive science (A. Clark, 1989).

The behavior of complex systems (e.g., human beings) can be predicted or explained from more than one interpretive stance. According to Dennett (1987), intentional stance explanations account for regularities in behavior (a) by attributing wants, goals, beliefs, and desires to behaving actors and (b) by assuming that the actors will act rationally in accordance with their beliefs, desires, etc. Such explanations are often termed folk psychological because of their dependence on the mentalistic vocabulary used in ordinary interaction. Dennett claims that complex behavior can also be described without reference to the propositional attitudes (i.e., thoughts, beliefs, etc.). Design stance explanations, for example, rely "on the notion of function . . . and design stance predictions are generated by assuming that each functional part will function properly" (1981, p. 4). The behavior of an alarm clock or thermostat can be predicted from knowledge of its design. Yet another possible mode of explanation can be carried out from the physical stance, where behavior is predicted from knowledge of the physical states and known dynamical relations between objects (Dennett, 1981). Planetary motion and the interactions of subatomic particles are typically explained from the physical stance.

Cognitive scientists have long recognized this 'levels of description' problem, but have often assumed that descriptions from the intentional stance could be reduced to the functional terms of the design stance and subsequently to the physical stance. This commitment is manifested in the standard cognitivist assumption that the beliefs and desires of folk psychology map neatly and isomorphically onto the computational architecture of the brain (for a statement of this position, see Fodor & Pylyshyn, 1988; for critiques, see Clark, 1989; Dennett, 1987).

Clark (1989) claims that this is a mistake. He argues that the rule-governed manipulation of syntactically structured symbolic representations fails to provide a computational substrate capable of supporting the rich and flexible behavior that is characteristic (say) of locally managed, situated action. Much of his argument centers on the notion that the ascription of thoughts, desires, etc. is made holistically. As Clark puts it:

> . . . we ascribe beliefs by throwing a kind of interpretive net over a whole body of behavior. And the mesh of that net is gauged to our particular interest in making sense of behavior. The knots (i.e., the particular ascriptions of beliefs and desires) need not correspond to any natural, projectible divisions in whatever underlying physical or computational structures make possible the behaviors concerned . . . . (1989, p. 49).

In other words, there may be no neat correspondence between the terms in our accounts of intelligent actions and the mechanisms that generate the actions themselves. But if the rules and symbols of classical cognitivism are not "in the head," what kinds of processes do support intelligent behavior? PDP-style processing may provide the necessary computational substrate.

To review, the planning model reflects the attempt to formalize everyday talk about mental states by internalizing the descriptive concepts of the intentional stance. The failure of the planning model to answer to the details of local management is deeply related to the failure of decontextualized, syntactically structured propositional attitudes as computational primitives. The flexible performance afforded by PDP systems seems like just what is needed to model the complexity and context sensitivity of locally managed, situated action. This is not to say that talk of plans or of beliefs and desires is useless. On the contrary, standard and PDP models serve complementary roles in psychological explanation.

Standard cognitivism is one basis for descriptive cognitive science. Descriptive cognitive science "attempts to give a formal theory or model of the structure of the abstract domain of thoughts, using the computer program as a tool or medium" (Clark, 1989, p. 153). PDP cognitivism is one basis for causal cognitive science. Causal cognitive science "attempts to give an account of the inner computational causes of intelligent behaviors that form the basis for the ascription of thoughts" (Clark, 1989, p. 154). Plan-based models of message production are descriptive in the sense just described. The purpose of this thesis is to construct a PDP model of message production that is causal in the sense just described. Such a model is intended to capture aspects of message behavior that are highlighted in local management conceptions of intelligent adaptation.

Summary

This chapter has introduced a number of oppositions: planning versus local management, symbolic versus PDP modelling, descriptive versus causal cognitive science, and formal versus content-based models of message production. These oppositions run throughout the whole thesis, and no effort was made to flesh them out fully in this introduction. The attitude adopted toward these contrasts does not force one to choose sides. Rather, the challenge is to point out where each perspective ought to be applied (A. Clark, 1989; Suchman, 1987). To forecast things slightly, formal, symbolic planning models are best seen as descriptive models of the rational accountability of interaction. Local management, a la conversation analysis, is a descriptive theory of situated activity as it actually unfolds. The PDP model I propose is a causal theory of the local organization of expressed content in message production.

Primary Research Questions

This thesis is an attempt to begin to answer five interrelated questions about message production and communicative action: (a) What is the cause or origin of functional diversity in message design? (b) What kinds of representations and processes underlie message production? (c) Is the nature of communicative action best captured in terms of rational planning or local management, or neither, or both? (d) Are plan-based accounts of message production empirically adequate? (e) Do parallel distributed processing models offer a credible alternative to plan-based models of message production?

Outline of Chapters

The thesis is developed in eight chapters. This introductory chapter has contrasted two models of adaptation commonly used to understand cognition and communicative action. Rational, plan-based models of cognition were compared to models of situated action and local management. Several theoretical distinctions were introduced, and primary research questions were stated.

Chapter 2 surveys comprehensive models of message production that attempt to describe the sequence of processes that bridge the gap between "intention and articulation" (Levelt, 1989). The planning model is described and its roots in philosophy, psychology, artificial intelligence, and cognitive science are examined. Chapter 3 is a critical evaluation of plan-based models. The chapter contends that plan-based models offer an explanation of message production that depends on the untenable assumption that there is a stable relationship between abstract communicative functions and decontextualized linguistic forms.

Chapter 4 describes a parallel distributed processing (PDP), model of situated message production. O'Keefe's (1988) theory of message production is recast in this PDP modelling framework. Chapter 5 describes the development and validation of a thought checklist procedure for thought elicitation. Chapter 6 describes the development and validation of a system for coding messages into a numeric vector representation. Chapter 7 describes a PDP model of the mapping between patterns of thoughts and patterns of message elements. Chapter 8 summarizes the findings, considers their limitations, and points to directions for future work.

CHAPTER TWO
COMPREHENSIVE MODELS OF LANGUAGE PRODUCTION

I'm an educated man, to get more sense within my head I plan.
-Doc Watson, "Born About Six Thousand Years Ago"

The last decade of research in communication and psycholinguistics has produced several comprehensive models of language production (see e.g., Bock, 1982; Garrett, 1980, 1988, 1990; Greene, 1984a; Herrmann, 1983; Levelt, 1989). The models differ significantly, but it is possible to point out commonalities. For instance, there is a consensus that a model of message production ought to characterize the cognitive processes that begin with having a goal or intention and end with the physical articulation of an expression in language.[1] It is also widely agreed that language production is not a one-shot deal; many steps are required to get from thoughts to articulated expression. At the same time there is no agreement as to just how many subprocesses intervene between intention and articulation, nor is it clear just what these processes are, what representations they operate on, nor how they interact.

Overview

The purpose of this chapter is to describe plan-based models of message production. However, message production is only one step in the whole language production process.[2] Before discussing individual message production models, it is important to get a picture of how the entire language production systems is thought to operate.

The chapter is divided into several sections. As an organizing framework, Levelt's (1989) basic blueprint for the speaker is described in the first section. In the second section, message production is considered as a computational problem, and some common cognitive representations and processes are described. The third section introduces the AI notions of planning and problem solving and shows how they have been used in models of message production. The fourth section focuses on the use of abstraction in plan-based models and points out that indexing of linguistic forms is the key theoretical move that allows abstract descriptions of actions (e.g., sequences of speech acts) to be transformed into concrete linguistic forms, and the final section is a summary.

Three Processes in Language Production

In what follows, the stages of processing in production are labelled according to the system laid out in Levelt's (1989) "Blueprint for the speaker" (see Figure 1). The choice of Levelt's terminology does not stand as an uncritical endorsement of his decomposition of the process of production. It is adopted (a) because it is a useful heuristic framework for organizing a complicated literature, and (b) because it represents the standard view in the field today. Levelt breaks the process of production down into three main subprocesses: (a) conceptualization, (b) formulation, and (c) articulation.

Conceptualization

Processing in this subcomponent encompasses all of the thinking that precedes speech. Both discourse and message-level planning decisions are made by this subcomponent, which also monitors ongoing overt speech (Clark & Clark, 1977; Levelt, 1989). This stage of

[1] For the present purposes, whether expression takes place in an oral or written medium will not be a concern.

[2] The standard view portrays language production as a multi-step process involving hierarchical structures at many levels of abstraction. Roughly speaking, message production has to do with deciding 'what to say'. Sentence production has to do with deciding 'how to say it'. Language production is used to refer to the whole process, from idea to expression.

processing, message-level planning in particular, is the central focus and concern of the present thesis. For the time being, only a sketch is given of the processing involved here. After briefly describing the processes of linguistic formulation and articulatory execution, we will return to meditate on many of the issues in message-level planning.

Most discussions of conceptualization begin by emphasizing the instrumentality of human language use, noting that people speak primarily in order to achieve their goals (Clark & Clark, 1977; Herrmann, 1983; Levelt, 1989; Schlesinger, 1977). Thus speaking begins with a goal, an idea to be expressed, an intention. This intention is a conceptual or semantic, rather than a linguistic object. In addition, intentions that lead to linguistic expressions are of a special sort: they are communicative intentions, intentions that require the addressee to recognize the speaker's intention to communicate (Bach & Harnish, 1979; Levelt, 1989). The root intentions in a plan are thought to be the output of complex processes of comprehension, inference and interpretation, and are thus thought to lie outside the scope of language production research.

Once a speaker has a goal, an intention, or an idea (the terms can be used somewhat interchangeably), the next decisions to be made concern discourse type. At this juncture, speakers must adapt their ideas to the particular discourse genre (Clark & Clark, 1977) or, more accurately, participation framework (Goffman, 1981) that is currently in place. One must obviously put into place a different plan to give a lecture than to have a conversation or to give testimony in court.

Macroplanning. Message-level planning is broken down into macroplanning and microplanning (van Dijk & Kintsch, 1983; Levelt, 1989). Macroplanning elaborates a high level goal (macrointention) into a structured sequence of subgoals. For each subgoal, macroplanning selects the propositional content as well as the illocutionary force of each speech act, and assigns thematic structure to each message. The decomposition of goals into subgoals is a rational problem solving strategy that has deep roots in the study of human action, and it is a central part of most computer models of planning (Miller, Galanter, & Pribram, 1960; Newell & Simon, 1972; Barr & Feigenbaum, 1981). Finally, macroplanning processes linearize the message elements, ordering them for expression (Levelt, 1989). The output of macroplanning is an ordered sequence of speech acts.

Microplanning. Microplanning takes an ordered sequence of speech acts and adds to each act information the formulator needs to make syntactic, lexical and phonological choices. At this stage, message elements are indexed as given or new (H.Clark & Clark, 1977) and are indexed for such features as naturalness and vividness (Osgood & Bock, 1977). Focus, emphasis and style are also encoded at this level (Schlesinger, 1977).

The final processes in microplanning assign topic, perspective and accessibility indices to messages and anticipate any language-specific requirements of the formulator (Levelt, 1989). Such language-specific requirements might include providing temporal information to the formulator of a language with a tense system. The product of microplanning is a preverbal message, a conceptual structure that contains all of the information needed by the formulator.

Formulation

Formulation processes transform nonlinguistic, preverbal messages into linguistic objects. In common sense terms, this is the set of processes that translates ideas into sentences or thought into speech (Deese, 1984). The large literature on sentence production deals primarily with issues of formulation (Garrett, 1975, 1980, 1988). This aspect of language production is by far the most studied and best understood.[3] There are a couple of

[3] Levelt's (1989) text devotes roughly 90 pages to conceptualization, but more than 250 to formulation.

reasons for this to be the case. On the one hand, this is the area where the powerful formalisms of modern linguistics can most easily be brought to bear. On the other hand, the lack of well-developed methods for message analysis and elicitation so often bemoaned in the sentence production literature (see, e.g., Clark & Clark, 1977; van Dijk & Kintsch, 1983; Osgood & Bock, 1977) forces interested researchers to use speech error data, which are conventionally taken as evidence of formulation processes (Garrett, 1975, 1988).

Like conceptualization, formulation is further broken down by Levelt (1989) into two sub-processes: grammatical and phonological encoding. Grammatical encoding accesses lemmas from the mental lexicon and builds syntactic structures using the preverbal message as a blueprint. A lemma is "an abstract symbol representing the word as a semantic/syntactic entity . . . a kind of pointer or address to the words form" (Dell & Juliano, 1990, p. 1). The output of grammatical encoding is a surface structure tree. Phonological encoding then retrieves the phonological forms (or lexemes) corresponding to the lemmas contained in the surface structure, and outputs a phonetic or articulatory plan for the corresponding surface structure.

Articulation

The articulator takes as input the articulatory or phonetic plan from the formulator and executes the muscle movements that produce overt speech. Though a fascinating subject in its own right, a detailed discussion of articulatory processes is beyond the scope of this thesis.

This section described three basic processes in language production: conceptualization, formulation, and articulation. The purpose was to provide a larger framework within which to discuss problems specifically related to message production. Henceforth, the processes of formulation and articulation will not be a concern. The remainder of this chapter deals exclusively with issues related to message-level production. A more detailed survey of knowledge representation and processing in message production is made, and specific proposals about the architecture of the message production system are reviewed.

Representation and Processing in Message-Level Production

Recall that message-level processing transforms an intention or goal (represented nonlinguistically) into an annotated sequence of speech acts. This complex product is a preverbal message, or simply a message. As Garrett points out, the message "is the proximal cause of sentence production: it is the real-time representation that controls the integration of sentence form. Diverse inferential processes underlie the construction of a message representation" (1988, p. 70). It is to these diverse inferential processes, and associated representations, that attention is now turned.

Knowledge Representation

Declarative knowledge. There is a strong consensus among production researchers that the knowledge underlying message production is propositional (Bock, 1982; Herrmann, 1983; Levelt, 1989; Schlesinger, 1977). This is not to say all knowledge takes the form of propositions. On the contrary, there is good evidence to suggest that the mind represents knowledge in a variety of different ways. Mental models (Johnson-Laird, 1983), images (Farah, 1988; Kosslyn & Pomerantz, 1977), and kinesthetic formats (Levelt, 1989) have all been offered at one time or another as candidate representational systems (for a review of the knowledge representation literature, see Rumelhart & Norman, 1988). Most language production researchers sidestep the main issues in knowledge representation, settling on a fairly standard propositional format.

Propositions are especially well suited to the task of representing preverbal messages, where reference, predication, and thematic (i.e., Fillmore, 1968) structuring must all be achieved. Among those who have advocated a propositional representation at this level, the

preferred notational system is a sort of predicate/argument notation using parentheses to set off predicates from arguments at different levels of abstraction. For example, Schlesinger (1977, p. 13) represents the sentence "Alice kicks poor Bill" as:

(AGENT-ACTION Alice, (GOAL-ACTION
(ATTRIBUTE poor, Bill), kick))

where AGENT-ACTION, GOAL-ACTION, and ATTRIBUTE are all predicates that take one or more arguments. Herrmann (1983), while much more hesitant to claim any psychological validity for propositions, uses a similar notational scheme. Herrmann (1983, p. 37) notates the conceptual representation of the sentence "Otto has a girlfriend named Irene" as follows:

HAVING (OTTO, GIRLFRIEND (IRENE))

Again, HAVING and GIRLFRIEND are predicates taking some specified number of arguments. In one last example, Levelt (1989, p. 81) notates the sentence "Cynthia saw that Joe put the key under the doormat" in the following manner:

SEE(CYNTHIA, PUT(JOE, KEY_1, UNDER($DOORMAT_3$)))

The only substantive difference here involves the use of subscripts to index unique tokens such as KEY_1 and $DOORMAT_3$.

Intuitively, one would suppose that virtually anything a person could know would have to be potentially available to the processes of production. Surveying the literature on this point, that is about how things look. Clark and Clark (1977, ch. 6) suggest that message production is a kind of problem solving that depends on several sources of knowledge: knowledge of the listener, of Gricean (1975) maxims for cooperation, of possible and impossible states of affairs, of social contexts, and of linguistic devices. Levelt (1989) organizes the different knowledge domains into three categories: discourse model, situation knowledge, and encyclopedic knowledge. The discourse model includes four important features: (a) a representation of assumed common ground shared by participants, (b) a record of one's own contributions, (c) a record of others' contributions, and (d) a record of information yet to be conveyed. In addition, the discourse model maintains pointers to that part of the discourse that is currently the topic or focus. Finally, the discourse model supports the current participation framework (Goffman, 1981).

The situation model captures non-propositional (especially spatial) aspects of discourse situations. Although such models seem to make important contributions to comprehension, their role in production is not well understood (Graesser & Bower, 1990; Morrow, Bower, & Greenspan, 1989; Morrow, Greenspan, & Bower, 1987; Perrig & Kintsch, 1985).

Reciting such a litany of knowledge domains does little to advance our understanding of the processes of message production, but it does point to the complexity of the process. In particular, one notices that message-level decisions are dependent on a wide variety of factors. That is, messages can be seen to covary with individuals' knowledge of facts, norms, situations, as well as with "external" situational factors. The standard view assumes that this covariation is the result of deliberate selection based on these factors (see, e.g., Brown & Levinson, 1987; Herrmann, 1983). This logic leads one to suppose that any situation feature or characteristic of speaker or hearer knowledge that covaries with some message structure must be the basis for selection and therefore must be explicitly available to message-level processes. Such a line of reasoning leads to serious trouble as the number of discernible message features grows (D. O'Keefe, 1987, 1990).

Procedural knowledge. The distinction between procedural knowledge (knowledge how) and declarative knowledge (knowledge that) has a checkered history in cognitive science (for a review, see Rumelhart & Norman, 1988), but has proved useful in characterizing knowledge about behavioral production. Bock (1982), Greene (1984a,

1984b), and Levelt (1989) each claim that procedural knowledge in language production is stored in the form of if-then rules called productions. Production rules are condition-action rules such as "If hungry then eat." A production system is an organized set of such rules operating under executive control. An action is carried out in a production system when its condition matches the current state of the world. Production systems have been used to model a wide variety of cognitive functions (Anderson, 1976, 1983; Newell, 1973, 1980; Newell & Simon, 1972).

Greene, for example, argues that procedural knowledge underlying communicative action is stored in a set of "procedural records" (1984a, p. 294) at various levels of abstraction. These records represent condition/action relations whose levels of activation and strength are a function of the current situation and past experience. They are indexed in procedural memory by the desired actions/outcomes and "any initial conditions which have proven relevant in past goal pursuits" (Greene, 1984a, p. 292).

Levelt (1989) chooses production systems for their generality (they can be used at each level of processing: conceptual, grammatical, lexical), and offers examples of the sorts of productions that might underlie message planning. In particular, productions that support the planning of speech acts have illocutionary intentions as their conditions and speech acts as their actions. Levelt develops a simple example as follows:

> IF the intention is to commit oneself to the truth of p, THEN assert *p*. Here *p* is some proposition the speaker wants to express as being the case, and the indicated procedure is to build an assertion of that proposition. The Conceptualizer and its message generator can be thought of as a structured system of such condition/action pairs. (1989, p. 10)

Detailed systems of production rules for language production have also been proposed by researchers in artificial intelligence, and these are explored below.

Plan-Based Models of Message Production

Plans and planning play a central role in cognitive theories of communication, especially in theories of message production (Berger, 1987, 1988; Berger & Bell, 1988; Dillard, 1990; Kellerman et al., 1989; Pavitt, 1991; Waldron, 1990). In most instances, theories of planning have been drawn from research traditions in cognitive psychology (Hayes-Roth & Hayes-Roth, 1979; Hobbs & Evans, 1980) computer science (Newell & Simon, 1972), and artificial intelligence (Fikes & Nilsson, 1971; Sacerdoti, 1974, 1977), where the planning concept has been refined by writing computer programs that model the planning process.

In contrast, communication researchers have not engaged in modelling, and an unfortunate consequence has been the uncritical acceptance or at least the lack of recognition of certain theoretical assumptions about planning mechanisms that would probably not be accepted if made explicit. Demonstrably false assumptions about form/function relations in particular are common to all plan-based models but have received little attention or comment from communication theorists who endorse plan-based approaches.

This section examines plan-based models of message production that have actually been implemented as computer programs. Implemented models are not superior to experimentally tested models; in fact both are necessary. However, it is in building computer implementations that all of a plan-based theory's assumptions must be specified. Once these assumptions are made explicit, they will be critiqued.

AI, Problem Solving, and Planning

Plan-based models of message production are founded on three distinct contributions from AI and cognitive psychology: (a) a means-ends analysis approach to problem solving (Newell, Shaw, & Simon, 1963; Newell & Simon, 1972), (b) a AI approach to representing

and processing plan knowledge (Fikes & Nilsson, 1971; Miller, Galanter, & Pribram, 1960; Sacerdoti, 1974, 1977); and (c) an analysis of utterances as speech acts (Bruce, 1980, 1986; Cohen & Perrault, 1979; Searle, 1969). Taken alone, each of these contributions is associated with a large literature. No attempt is made here at a systematic survey of these literatures. Instead, each basic idea is explained, and an effort is made to highlight those features of plan-based approaches that will be relevant in subsequent discussion.

Problem solving. Newell and Simon's (1972) seminal contribution to cognitive psychology and AI was to produce a detailed model of problem solving, a model so detailed it was the basis for computer simulations of human problem solving performance. The General Problem Solver (GPS), as their simulation was called, is widely recognized as the first computer simulation of human cognition. Though hindsight reveals GPS to be limited in many respects and not truly a 'general' problem solver (Boden, 1988), its status as an exemplar is beyond question. In many important respects, subsequent advances in problem solving and planning have been elaborate footnotes to GPS.

For Newell and Simon, problem solving involves a problem space, operators for moving around the problem space, and a control strategy for choosing operators (Barr & Feigenbaum, 1981). The problem space is a symbolic representation of a problem that encodes the current state and the goal state. Operators are abstract actions that map one problem state onto another. Using means-ends analysis as a control strategy, operators are chosen whose application to the problem reduces the difference between the current state and the goal. Thus, problem solving proceeds by applying a sequence of operators that trace a path through problem space from current state to goal state, and the sequence of operators that leads from current state to goal is the problem solution.

In order to solve a problem, one must first develop a representation of the problem space. According to Langley, Simon, Bradshaw, and Zytgow (1987, p. 8):

> The brain solves problems by creating a symbolic representation of the problem (called the problem space) that is capable of expressing initial, intermediate, and final problem situations as well as the whole range of concepts employed in the solution process, and using the operators that are contained in the definition of the problem space to modify the symbol structures that describe the situation (thereby conducting a mental search for a solution through the problem space).

The set of available operators in a problem space constitutes the means by which the problem solver can alter the problem situation. Different problems carry with them different sets of operators. For instance, a chess-playing program is endowed with operators to move the various pieces around the board, while a theorem prover might have operators to remove the existential quantifiers from all the axioms in a list. In many models of message production, speech acts are conceived of as operators that change one intentional state to another (Cohen & Perrault, 1979).

Search through the problem space is guided both by a representation of the goal (or some means for measuring the distance to the goal) and by heuristics or rules of thumb which are provided by the programmer (Simon, 1977). Heuristics commonly take the form of production rules such as: "If the first step in a path leads you closer to the goal state, pursue that path further" or, using a familiar example, "If the goal is to commit to the truth of *p*, then assert *p*" (Levelt, 1989, p. 10). Intelligent problem-solving behavior depends crucially on the ability to develop powerful heuristics and to understand the triggering conditions and effects of all the operators in a given problem domain.

In the absence of strong, domain-specific heuristics, weaker and more general methods are used to limit the search for appropriate operators. Means-ends analysis is the term given

to a class of weak heuristics that has proved useful in a variety of problem domains. Means-ends analysis "refers to the process of comparing what is given or known to what is desired and, on the basis of this comparison, selecting a 'reasonable' thing to do next" (Ernst, 1987, p. 578). In slightly more technical terms, means-ends analysis describes a situation "in which a test heuristic detects specific differences between the current situation and the goal and evokes specific move operators that are relevant to removing just those kinds of differences that have been detected" (Langley et al., 1987, p. 12). The effectiveness of means-ends analysis depends upon the problem solver's ability to recognize and prioritize salient differences between current and goal states and on the accessibility of the operators necessary for reducing such differences once identified (Simon, 1977).

To sum up: A problem space represents all conceivable states of a problem. Important states include the current state and the goal state. Differences between current and goal states are reduced by the application of relevant operators. Operators are abstract actions that map one problem state onto another and are typically indexed by the differences they reduce. Finally, heuristics, or rules-of-thumb, guide the search for an appropriate sequence of operators.[4] The next section describes AI planning techniques based on these principles.

Planning. The notion of planning, though part of common sense and philosophical discussions of rational action since ancient times, was reintroduced to modern behavioral science by the work of Miller, Galanter, and Pribram (1960). Dissatisfied with behaviorism, Miller et al. wanted to "describe how actions are controlled by an organism's internal representation of its universe" (1960, p. 12). They were concerned especially to account for actions that are organized hierarchically and sequentially. As such, their oft-repeated definition describes a plan as "any hierarchical process in the organism that can control the order in which a sequence of actions is to be performed" (Miller et al., 1960, p. 16). The concept of planning has become a cornerstone of cognitive theorizing.

Speech Acts and Plans

The idea of combining speech act theory and plan-based AI into a single conceptual framework was first explored in a short but influential paper by Bruce (1986, originally published in 1975) entitled "Generation as social action." In this paper, Bruce's main point was that "generation needs to be understood as an action in a social context" (1986, p. 420). If words and sentences are taken at face value, Bruce argued, it is difficult to get a grip on the conceptual relations that make discourses coherent. His contribution was to show that the relations between utterances are seen as coherent because they are interpreted as component actions in rationally designed plans. According to Bruce (1986), "The production of a natural language utterance is best understood as an action which alters the state of the world, rather than as a mapping from a given meaning representation into a surface structure" (p. 419). In essence, a pragmatic approach to production (what Bruce calls generation) is endorsed over a semantic alternative.

In retrospect, combining speech act theory (Searle, 1969) and plan-based AI seems to be an obvious research strategy. In Searle's (1969) account of felicity conditions one has a virtually ready-made set of preconditions for STRIPS-like operators.[5] In addition, our

[4] More thorough discussions of problem solving are available in Barr and Feigenbaum (1981), Newell and Simon (1972), and Ernst (1987).

[5] STRIPS was the acronym given to an early planning formalism (Fikes & Nilsson, 1971). An operator in STRIPS is defined by three separate bits of information: (a) its preconditions, (b) its add list, and (c) its delete list. The preconditions are those conditions that must be true of the current state in order for the operator to be applied. The add list specifies what new formulas are added to the current state as a result of the operator's application. The delete list specifies which formulas are deleted from the current state as a result of the operator's

everyday vocabulary of intentional attribution (Dennett, 1987) captures information about the intended and actual effects of speech acts. All that is left is to see speech acts as operators in a rationally and conventionally defined problem space. Of course, by the time of Bruce's writing, Austin (1962) and Searle (1969) had already developed, in speech act theory, a detailed philosophical framework for interpreting utterances as actions.

What made Bruce's contribution truly seminal was the potentially fruitful linkage he saw between speech act theory and plan-based approaches in AI and cognitive science (Grosz, Sparck-Jones, & Webber, 1986). The commitment to plan-like structures as an integrative framework is made obvious by the manner in which social actions are defined. In addition to being about beliefs, desires, wants, etc., a social action "is best thought of as a set of operations to be performed on a belief system" (Bruce, 1986, p. 419). It is a short distance from operations on belief states to operators in a problem space.

The insight that speech acts could be modelled as operators in a planning system, though not explicitly stated until later (Cohen & Perrault, 1979), is reflected in the proposed definition of actions. "The full representation of an action," Bruce notes, "seems to require a representation of both its actual and its intended effects, its actual and its assumed preconditions" (1986, p. 419). Speech acts, just as operators in GPS or STRIPS, are defined in terms of their preconditions and their effects.

By crafting a plausible linkage between speech act theory and the AI notion of planning, Bruce opened up a whole new space for work on language production and comprehension. In doing so he gave birth to the problematic that organizes the study of language production still. This problematic was subsequently taken up by several researchers in AI (Appelt, 1985; Cohen & Perrault, 1979) whose work is the subject of the next section.

Speech Act Models of Message Planning and Production

Rather than attempting an exhaustive review of computer models of language production (see Kempen, 1987; McDonald, 1987; McDonald & Bolc, 1988; McKeown & Swartout, 1988), this section examines one well known model of conceptualization (Cohen & Perrault, 1979) and one system that models both conceptualization and formulation (Appelt, 1985). Cohen and Perrault's (1979) model was the first to implement a planning system using speech acts as operators in a problem space defined by agents' wants, beliefs, and intentions. This system planned sequences of acts to achieve goals, but did not map the acts onto utterances. Appelt (1982, 1986) extended this work in the domain of referring expressions and presented one way of mapping acts onto utterances. Each model is considered in turn.

Plan-based models of speech acts. The sections leading up to this one have explored the notions of planning, problem solving and speech acts. Cohen & Perrault (1979) put all of the pieces together in a computer program that modelled speech acts as operators in a planning system. Adopting a STRIPS-style planning representation (Fikes & Nilsson, 1971), operators were defined in terms of preconditions, bodies and effects. The preconditions must be true of the current state in order to apply an operator. Effects were descriptions of the changes in state brought about by the application of an operator. The body of an operator specified the means by which an operator had its effects.

The authors summarized the central thrust of their work in the following statement:

We hypothesize that people maintain, as part of their models of the world, symbolic

---

application. Together, the add and delete lists define the effects associated with the application of a given operator. An operator is chosen if its add list contains a condition that is present in the goal but not in the current state. An operator may also be chosen if its delete list contains a condition that is not in the goal but is in the current state (see also Barr and Feigenbaum, 1981; Ernst 1987).

> descriptions of the world models of other people. Our plan-based approach will regard speech acts as operators whose effects are primarily on the models that speakers and hearers maintain of each other. (Cohen & Perrault, 1979, p. 179)

Cohen and Perrault's (1979) modelling experiments demonstrated that aspects of message-level planning could be usefully modelled using speech acts as operators in a standard planning scheme.

Planning referring expressions. Appelt's (1982, 1986) planning system, KAMP (Knowledge And Modalities Planner), was designed to produce English referring expressions to assist people in the assembly of complex mechanical objects. It is a plan-based model of language production in the same tradition as Bruce (1975) and Cohen and Perrault (1979). Two important advances distinguish Appelt's KAMP from Cohen & Perrault's model: (a) Appelt (1986) uses a NOAH-style planning system (Sacerdoti, 1977)[6] as the basis for his model, whereas Cohen and Perrault used a version of STRIPS (Fikes & Nilsson, 1971); and (b) Appelt's system produces actual English sentences; whereas Cohen and Perrault only produced sequences of speech acts.

The planning of referring expressions requires the planner to reason extensively about the knowledge and intentions of both speaker and hearer. For this reasoning to take place, there must be a formalism for representing knowledge about intentions, wants, beliefs, and actions. Appelt (1986) used Moore's (1980) possible-worlds semantics of propositional attitudes to represent this knowledge. Illocutionary acts such as requesting and informing were defined as functions on possible worlds. As Appelt notes (1986, p. 504), "The semantics of action-terms . . . are axiomatized as relations on possible worlds, where actions map one world onto another, with possible worlds playing a role similar to the familiar states of a conventional logic of action." The use of such a sophisticated logical formalism can be daunting, but the image of production projected here is essentially the same as those that have come before. Speech acts are still treated as problem solving operators that map one problem state onto another.

Appelt's (1986) unique contribution involved mapping illocutionary acts onto surface speech acts (i.e., concrete words and phrases). The pervasiveness of indirectness and other forms of nonliteral meaning makes this a notoriously difficult problem (see, e.g., Brown & Levinson, 1987; Goffman, 1981; Levinson, 1979a, 1983; Searle, 1975), and Appelt is aware of the difficulty:

> The same surface speech act can realize different illocutionary acts, depending on the context in which it is performed, and it is possible for a single surface speech act to realize several illocutionary acts simultaneously. Because the effects of a given surface

---

[6] NOAH (Nets Of Action Hierarchies), was developed by Sacerdoti (1977) as part of the computer-based consultant project at Stanford Research Institute. NOAH differs from GPS, STRIPS, and ABSTRIPS in that it does not use simple means-ends analysis to select operators. Instead it relies on (a) a procedural net, (b) detailed domain-specific knowledge compiled as procedural rules (a.k.a. "SOUP" functions), and (c) critics.

NOAH begins by creating a single node to stand for the highest level goal. That goal is then expanded into subgoal nodes by the application of SOUP functions. SOUP functions represent procedural knowledge about how to solve specific problems. They are compiled procedures that may either expand a node into subgoals and suggest further SOUP functions to apply. Otherwise, they may simply indicate the appropriate primitive operator to apply, without suggesting further expansion. A node is expanded no further when it represents a subgoal that can be accomplished by the application of a primitive operator.
SOUP functions are indexed by the subgoals they apply to.

> speech act depend on what illocutionary act it is being used to realize, it is impossible to describe its effects directly. (1986, p. 508)

A successful plan-based approach to mapping illocutions onto surface speech acts in the face of these difficulties would have been a triumph for plan-based theories of discourse generally.

Unfortunately, Appelt chose not to tackle indirectness "because indirect speech acts are frequently planned to satisfy multiple goals, often along social dimensions and others that are very difficult to formalize" (Appelt, 1986, p. 509). Instead only direct speech acts were planned and the form/force issue was dealt with on an ad hoc basis by statically associating specific sentence types (e.g., declarative, imperative, interrogative) with specific illocutionary forces (e.g., informs, commands, requests). In practical terms, this is a solution that works, but as the basis for a theory of message production it does not.

Summary of plan-based models. Message production is treated as a special case of problem solving. Planning messages is essentially a two step process. First, goals provided by the modeller are transformed into abstract acts. Second, abstract acts are transformed into concrete linguistic forms. Both transformations require a kind of indexing. Abstract acts are indexed by the effects they have on the world, and linguistic forms are indexed by the abstract acts they instantiate. Any system in which the planning mechanism operates on abstractions will require an indexing system to instantiate the abstractions. The following section examines the role of functional abstraction in plan-based models of message production generally.

Functional Abstraction in Plan-Based Models

Both of the plan-based models considered so far have in common a commitment not just to speech act theory but to a kind of stable form/function (or form/force) distinction.[7] Using some more or less formalized theory of action, these models begin by computing abstract functional descriptions of utterances. These functional descriptions are then used to select linguistic forms to instantiate at progressively lower levels of abstraction. The selection is possible because the linguistic forms are indexed by their functional descriptions. This basic framework is what I call a functional indexing scheme.

Functional Indexing Schemes

A functional indexing scheme (FIS) is a system for storing and retrieving linguistic forms according to some set of abstract functional descriptors. The actual indices in a given FIS may vary. Appelt (1986) uses speech act categories to index syntactic frames, while others (see esp. Hovy, 1990) have developed more elaborate sets of pragmatic and rhetorical features for indexing linguistic forms. Regardless of the indices chosen, the basic idea is always the same: Linguistic forms at many levels of abstraction are indexed by their functional (or pragmatic or rhetorical or whatever kind of) description. Planning is a process of building up an abstract representation of an utterance in terms of the descriptors in the chosen FIS. The abstract representation is then transformed into an actual utterance by using the functional indices to select the appropriate linguistic forms. It is vital to recognize that any system which employs abstract functional descriptions in planning utterances must use some sort of indexing scheme to move between levels of abstraction.

Example of a functional indexing scheme. The complexity of computer models of planning and problem solving may partially mask the underlying reliance on functional indexing. Schlesinger's (1977) theory of language production is instructive because it provides an explicit theoretical statement endorsing a functional indexing solution to the problem of instantiating abstract functional descriptions. Schlesinger realizes that abstract,

[7] The distinction is apparent in the separation of propositional content from illocutionary force (see e.g., Levelt, 1989; Searle, 1969).

conceptual representations may not uniquely determine linguistic forms. For example, the nonlinguistic, semantic concept 'overweight' might be instantiated as overweight, fat, obese, chubby, heavy, large, Rubenesque, etc. How, then, does the message production system instantiate a semantic concept which does not specify a unique word? The solution goes like this: If semantic considerations alone fail to specify unique linguistic forms, then enrich the representation of linguistic forms with functional information. Forms will then be indexed not only by their meaning, but by their communicative function as well. Alternative words may mean the same thing, but, the argument goes, they have different functions. In the course of producing a message, where semantic considerations fail to select unique forms, functional descriptors serve as tie-breakers.

Word choice is not the only operation under the control of communicative considerations (functional indices). Alternative syntactic realizations of illocutionary force are also indexed by the communicative considerations (functions) they satisfy. As Schlesinger notes:

> When making a request, one has various syntactic forms at one's disposal: 'imperatives (Give me my book back), declaratives (I would rather you give me my book back), and interrogatives (May I have my book back?). Similarly, there are various ways of expressing a question . . . we conclude that there is a communicative consideration of polite requesting, another one of indirectly requesting, and so on . . . it seems reasonable to postulate that to each communicative consideration, or perhaps to each combination of communicative considerations, corresponds a shunting operation influencing the form of the utterance. (1977, p. 70)

Here as elsewhere (see e.g., Appelt, 1986) it is suggested that the one-to-many mapping of illocutionary force to sentence type can be managed by introducing functional indices that specify which sentence types realize which illocutionary forces. Schlesinger (1977) offered this as a testable hypothesis about language production. The evidence bearing on this hypothesis is examined in the following chapter.

Summary

Language production is currently conceived as a multi-step process whereby hierarchically organized, nonlinguistic conceptual representations are constructed and then instantiated at progressively lower levels of abstraction (Dell & Juliano, 1990; Levelt, 1989). The initial construction of a nonlinguistic conceptual representation is driven by an image of the current state and the goal. The output of this construction process is an abstract functional description of the linguistic form(s) whose expression will transform the current state into the goal. Once an abstract description is constructed, language production consists in translating it into concrete linguistic forms.

In order to accomplish this translation, linguistic forms at each level of abstraction must be indexed in terms of the description used at the highest adjacent level. At the highest level is a goal or set of differences between current state and goal state. At the next level, acts are indexed by the differences they reduce. One level down, syntactic frames are indexed by the acts they instantiate. Moving still lower down the abstraction hierarchy, lemmas are indexed by the syntactic categories to which they belong; morphological and phonological forms are indexed by the lemmas they are associated with, and so on until the process bottoms out.[8] This kind of indexing and the concomitant reliance on conceptual and functional abstractions are definitive of current models of language production. The next chapter expands the description of this modelling framework and raises several questions about its empirical validity.

---

[8] This discussion is necessarily simplified. For a fuller tratment see Levelt (1989, esp. chapters 5 and 6).

CHAPTER THREE
A CRITIQUE OF TYPIFICATION IN COMMUNICATION THEORY

The best-laid schemes o' mice and men
Gang aft agley,
An' lea' us naught but grief an' pain,
for promis'd joy!
-Robert Burns, "To a Mouse"

The most insidious characteristic of plan-based models of message production is that they appear to work. Plan-based modellers have computer programs that seem to stand as existence proofs of the adequacy of this approach (see, e.g., Appelt, 1986; Hovy, 1990). What I try to demonstrate in this chapter is that the plan-based models, while perhaps adequate for practical applications, are woefully inadequate from the standpoint of communication theory.

The entire argument is based on a critical analysis of three common theoretical moves in language production research. Figure 2 is a schematic representation of these moves. Basically, models and theories of language production rely on typifications of situations, typifications of messages, and stable associations between the two sets of types.[9] The first mapping in Figure 2 represents the typification of situations. It is a common practice in communication theory to index situations by the communicative functions they involve. Thus the literature contains references to compliance-gaining situations, comforting situations, regulative situations, persuasive situations, etc. (Dillard, 1990; O'Keefe & Delia, 1988; Seibold, Cantrill, & Meyers, 1985). Alternatively, situations are indexed by the types of goals or differences-to-reduce that they embody (Levelt, 1989; Newell & Simon, 1972).

In a standard model (Appelt, 1986; Levelt, 1989; O'Keefe, 1990a), types of situations are then statically associated with typifications of messages as in the second mapping in Figure 2. The third mapping in Figure 2 represents the typification of message contents. The two most common typifications of message contents are speech acts (Appelt, 1986) and strategies (Seibold, Cantrill, & Meyers, 1985), though other typifications exist as well (see, e.g., the pragmatic and rhetorical features of Hovy, 1990). Finally, the fourth mapping in Figure 2 represents a potential direct mapping between thoughts and messages. This mapping is not a familiar part of the literature, but its significance is explored at length in this thesis.

This chapter argues against each of the first three moves on the following grounds: (a) Communication is a situated activity, and therefore it is a mistake to index situations by the types of goals or functions they embody; (b) Decontextualized linguistic forms have no functional significance, and therefore, it is a mistake to index forms by the functions they achieve or the acts they instantiate; and (c) Communicative means-ends relations are not stable across individuals or over the life-span of development, and therefore it is a mistake to reify a single set of means-ends relations.

A case for the first point about indexing situations by the type of goals or communicative functions they embody is made here for the first time, although it reflects the influence of several different authors (Lave, 1988; Suchman, 1987). Second, to demonstrate that decontextualized linguistic forms should not be indexed by the acts they instantiate, I extend Levinson's (1979a, 1979b, 1983) critical analysis of speech act theory to cover

[9] Of course, there is also a reliance on functional indexing schemes to instantiate the abstractions.

functional descriptions generally (D. O'Keefe, 1987). Third, in order to demonstrate the instability of means-ends relations, I draw on evidence about message diversity (Dillard, 1990; Lambert & O'Keefe, 1990; O'Keefe, 1988; Seibold, Cantrill, & Meyers, 1985).

The Typification of Situations

For the purpose of modelling message production, situations are not treated in and of themselves. Instead, each new situation is treated as an instance of some general type of situation, and production decisions are based on the type not the instance. In communication theory, this amounts to seeing situations in terms of comforting, negotiation, etc. and making production decisions accordingly. In cognitive science, it amounts to seeing situations in terms of the types of goals they embody. Plan-based models in particular treat goals as differences between states of the world (Newell & Simon, 1972). If some situation requires state S_1 to be transformed into state S_2, it is identified as involving difference-type D_1. But a different act might transform S_3 into S_4 and still be an instance of difference-type D_1. Notice that two distinct paths through state space, two distinct differences, can be instances of the same type.

Recall now that message production is a kind of problem solving where the object is to produce an utterance that moves me from the current intentional state to the goal state. The crucial indexing information needed in the selection of an appropriate utterance is not the type information but the token information. What is needed is a function S() that takes my current intentional state (I_c) and my goal state (I_g) and returns the utterance or sequence of utterances ($U_{c,g}$) that reduces the difference between them:

(2) $S(I_c, I_g) = U_{c,g}$

This would seem to be the most direct approach (see the fourth mapping in Figure 2), but most plan-based approaches insert another step. Rather than mapping directly from the current and the goal state to an utterance, the difference between the current and goal state is first classified as being some specific type of difference (see Figure 2). The current and goal states are first fed through the classification function:

(3) $C(I_c, I_g) = D$

The equation in (3) takes the current state and the goal state and outputs a type of difference (perhaps a set of differences). The difference-type can then be used to select an operator since operators are indexed by the types of differences they reduce. However, it is the detailed difference information itself that is crucial. An act in a context is associated with a distinct path through problem space. Similar paths through problem space get the same functional labels. But since many spatially distinct paths get the same label, the label is not of much use when trying to retrieve a particular path.[10] Vital information is lost in the lost in the classification of specific differences into types.

I want to claim that using abstract difference-types to index operators leads to brittle system performance, and much of this brittleness is due to a lack of appreciation of context-specificity of desired effects (goals, differences, communicative functions, etc.). To say goals are context specific is to suggest that people reason with the goals at a low level of

[10] Imagine two vectors: (1, 1) and (1, -1). They both have the same length (= 1.414), but they are oriented differently. Using only functional labels as indices is like attending to length while ignoring orientation. Note especially that orientation can not be recovered from length alone.

abstraction. Abstract descriptions of goals or differences are decontextualized and thus lack the detailed information about the nature of a given goal that is needed for operator selection.

In sum, vital information is lost in in the mapping of particular differences into abstract-types. Operators (e.g., message contents) must instead be indexed by the differences under a richer description so that context-sensitive retrieval of operators is possible. Seen in this light, the typification of situations is troublesome, since it involves executing an information-depleting many-to-one mapping that must eventually be inverted. It seems more sensible to assume that such a typification does not occur, and that people rely instead on a richer, situation-specific representation of current and goal state.

Just as situations are treated as instances of situation types, messages are treated as instances of act- or strategy- types. The next section describes the problems associated with the typification of messages, focusing first on speech act theory and then on the practice of message typification more generally.

The Typification of Messages

In an influential paper, Levinson (1979a) argues that speech act theory cannot be used as a general framework for modelling the coherence of dialogues. A primary target of his argument was the claim that one could specify a kind of grammar of dialogue using speech act categories as primitive units over which to define sequencing rules. Such attempts, according to Levinson, rest on the presumption that the following constraints hold:

> (i) There are unit acts - speech acts - that are performed in speaking.
>
> (ii) Utterances are segmentable into unit parts - utterance units -each of which corresponds to a unit act.
>
> (iii) There is a specifiable conventional procedure that will assign unit acts to utterance units; in other words, there is a function whose domain is the set of possible utterance units and whose range is the set of possible speech acts.
>
> (iv) Sequences of acts are regulated by a set of conventional sequencing rules stated over speech act-types. (1979a, p. 95)

There is ample evidence (presented below), however, that these constraints do not hold.

Whereas Levinson argued against the possibility of building a speech act-based grammar of dialogue, the same counterexamples bode poorly for plan-based models of language production. To the extent that standard plan-based models of production rely on (i) - (iii) above, Levinson's arguments ought to apply equally well to them.[11] First, the dependency of speech act models of production on these premises is displayed, and then Levinson's (1979a) reasons for questioning the validity of the premises are reviewed.

Standard Models Depend on the Premises

A commitment to the psychological reality of speech acts is obvious in plan-based models (Appelt, 1986; Bruce, 1986; Cohen & Perrault, 1979; Levelt, 1989). Previous sections have described how modellers capitalized on the fit between speech act theory and standard AI conceptions of planning and problem solving. Speech acts are modelled as operators in a planning system. Felicity conditions are transformed into the preconditions on these operators, and the conventional effects of doing particular acts are codified in the so-called add- and delete-lists of the operators (Fikes & Nilsson, 1971). Levinson claims that the level of illocutionary force is attractive "because on this deeper level rules of sequencing will be straight-forwardly describable" (1979a, p. 96). A related benefit makes act-level

[11] Levinson (1979a) attacks (iv) also, but that premise is not relevant for the present purposes. It refers to sequencing rules, but the plan-based models I have been reviewing do not rely on fixed sequencing rules. In these models, sequences of acts are selected based on their conventional effects in the context of a planning system.

descriptions attractive to plan-based modellers: it is easier (within a standard planning framework) to generate abstract descriptions and then instantiate them than it is to generate concrete message contents directly. The vocabulary of acts provided by speech act theory, once translated into a STRIPS- or NOAH-like formalism, makes it possible to plan sequences of actions apart from their realizations as utterances. But this convenience in planning creates enormous difficulties when it comes to realizing the acts as actual utterances.

Premise (ii) refers to the existence of independently specifiable utterance units that map to unit acts. This premise is irrelevant to modellers like Cohen & Perrault (1979) because they stopped short of trying to produce utterances. But it is an absolute necessity to those who try to map acts onto utterances (Appelt, 1986; Hovy, 1990). Such proposals rely on this premise because planning proceeds in a top down manner, first at the level of acts, and then at the level of utterances. In Appelt's (1986) KAMP system, for example, the appropriate illocutionary force is chosen, then a syntactic frame for realizing the illocutionary force is selected. This procedure must make reference to a database of linguistic forms (e.g., syntactic frames) wherein each form is indexed by the function it serves. But this indexing scheme presumes that there is a privileged grain size or level of structure (perhaps the sentence) that ought to be indexed by illocutionary force.

As for the third premise above, which asserts the existence of a specifiable procedure for mapping utterances into act-type categories, speech act models of language production rely on an even more stringent presumption. Such models need to map from acts to utterances. That is, speech act models of production presume not only that an utterance-to-act procedure exists but that its inverse exists (and is specifiable) as well.

Validity of the Premises

For each of the three premises that underlies attempts to build speech act based models of language behavior, Levinson (1979a) provides an analysis of counterexamples that raises grave doubts about their validity. The first premise asserts that there are unit acts performed in producing utterance units. This premise runs into trouble on two fronts. First, utterances in context may perform more than one act at a time. Indirect requests, as noted by Levinson (1979a) are just one example of this kind of double- or triple-duty utterance. The following example comes from an old but not so subtle friend of mine, who upon entering my house and looking in the refrigerator would innocently inquire:

(4) Is that beer I see in the refrigerator?

Clearly the utterance in (4) is both a question "Is that beer?" and a request "May I have a beer?" Double duty utterances also appear in other contexts as well. Utterances like the first contribution in (5), from Levinson (1979a), count as both questions and pre-offers:

(5) A: What are you doing tonight?
B: Nothing, why?
C: I was thinking of going to a movie, wanna come?

Nonlinguistic contextual factors can also give questions the force of complaints or answers the force of requests as in (6), overheard while working at on-campus registration:

(6) A: Can I help you?
B: Speech Com. 230?
A: How 'bout Speech Com. 335?

Every contribution in this short dialogue performs at least two functions. The registrar (A) opens the encounter. The student (B) wants to add a course (Speech Com. 230), but it is closed, so the registrar suggests an alternative (Speech Com. 335). A's first contribution is both a question and an offer. B's contribution is both an answer to A's question and a specific request. Finally, A produces a reply to B's request in the form of a question that also functions as a counter-offer. Many more examples of this and other types could be produced,

but these suffice as an illustration of the general point that utterances do not always correspond to unit acts. If in fact (i) does not hold, and utterances often function in many ways simultaneously, then: (a) it will be difficult to index utterances by their associated functions (since these will be many), and (b) it will be difficult to plan, since utterances designed to perform one act might perform many other unintended and potentially counterproductive acts simultaneously.

The second problem for (i) is that utterances can have their primary force not as act-types (requests, promises, offers, etc.) but as evidence of an underlying plan or motivational structure. Once an underlying plan is recognized, a hearer may respond not just to the previous utterance but to any issue relevant to the speaker's underlying plan. As Levinson (1979a, p. 112) notes, "given an utterance which is merely the first in a sequence predicted by a hierarchical structure of goals, one is free to respond to one of these higher level goals." This is a problem for (iv) as well as for (i). Planners must not only anticipate the multifunctionality of utterances; they must also anticipate whether an utterance suggests an underlying (again perhaps unintentional) plan structure. While systems capable of sophisticated plan inference and recognition (e.g., Allen, 1983; Allen & Perrault, 1980) might be able to accomplish these feats, it is not clear how existing systems will surmount the difficulties raised by the invalidity of premises (ii) and (iii) described below.

Levinson makes several points about the invalidity of (ii). Only one of these points is developed here. Speech act models of language production assume that there is a privileged utterance unit (typically the sentence) indexed by illocutionary force. This assumption is untenable primarily because utterance units are often not specifiable in advance of functional classification. According to Levinson (1979a, p. 103), "Which unit is relevant for speech act assignment cannot be determined in advance, for utterance units seem to be identified on functional grounds." Consider the following example (from Levinson, 1979a):

(7) A: Johnny, did <u>you</u> smear Susie's face with the paint?
B: (silence)

What is relevant about this example is that Johnny's response (by its official absence) would be interpreted as admission of guilt. The general point is that any unit, even silence, can carry illocutionary force. Further, just what the appropriate unit is depends on how an utterance is interpreted. But if proper units are not specifiable out of context or in advance of their interpretation, then schemes such as Appelt's (1986) that index sentence types by illocutionary forces are at best misguided.

Premise (iii) is probably the most significant for speech act models. It states that there is some specifiable utterance classification procedure or function that takes utterances as input and produces act-type classifications as outputs:

(8) $F(U_i) \rightarrow A_i$

where U_i is some utterance and A_i is the act-type category to which it belongs. There are many reasons to doubt that such a function could be specified. First off, as Levinson notes "the domain of the function is defined in part in relation to the range" (1979a, p. 103). In addition, the so-called double duty utterances described above would need to be mapped not to one act category but to many (i.e., the classification procedure might not be a function at all but rather a relation).

But the most intractable problem for this project is indirectness. It is one thing to classify utterances based on the presence of explicit performative verbs or based on some fixed sentence type to illocutionary force mapping,[12] but it is quite another to assign forces

to indirect utterances. One problem with indirect utterances is that their proper classification is not simply a function of the utterances themselves but may also be a function of context. This observation leads some to believe any viable classifier must be a function that maps utterance-context pairs into act categories (Gazdar, 1981; Levinson, 1979a, 1979b, 1983).[13]

Even if one were to grant that the utterance-to-act mapping could be specified, speech act models of language production would still not be saved. This is because models of production need the inverse mapping; that is, they need to be able to map acts onto utterances. But "for any given functional role there are likely to be an indefinitely large set of utterance units that will perform that function" (Levinson, 1979, p. 103). This is a fancy way of saying that the mapping from acts to utterances is one-to-many.[14] So, even if speech act modellers could specify the classification function (and no one has), they would still have to specify the inverse. And, having specified the inverse, they would have to offer a principled basis for choosing among the many alternatives that the inverse model would produce. Such a solution is not impossible in principle, but classical planning approaches seem ill-suited to the task.[15]

In summary, Levinson (1979a, 1979b, 1983) draws attention to several troubling features of speech act theory: First, a single utterance unit may perform many acts simultaneously. Second, utterance units are defined in part in relation to the acts they perform, so there can be no privileged, indexable utterance unit prior to some functional classification. Third, there is no empirically adequate procedure for mapping utterances into speech act categories, and even if there were, speech act modellers of language production would need to invert the mapping to support act-to-utterance reasoning.

One important implication of these criticisms of speech act theory pertains to the the viability of various functional indexing schemes that have been proposed to support the selection of concrete utterance content in production (see, e.g., Appelt, 1986; Hovy, 1990; Schlesinger, 1977). The work of Levinson (1979a, 1979b, 1983) and others (Gazdar, 1981; D. O'Keefe, 1987) undermines these schemes by highlighting the instability and

[12] This strategy also fails. Appelt (1986), in trying to avoid the complications brought on by indirectness, only works with direct utterances and uses a fixed mapping to get from act-type to sentence type. This kind of form-force mapping has no empirical justification whatsoever. As Levinson (1983, p. 265) notes, "On the face of it, what people *do* with sentences seems quite unrestricted by the surface form (i.e., sentence-type) of the sentences uttered." Add to this the observation that "*most* usages are indirect" (Levinson, 1983, p. 264) and one has a strategy "in which case assignment is a relatively simple affair but is largely irrelevant to how conversation proceeds" (Levinson, 1979a, p, 106).

[13] Even this approach seems problematic. Once we try to understand the functions of utterances in context, Levinson argues, "it soon becomes clear that the contextual sources that give rise to the assignment of function or purpose are of such complexity and of such interest in their own right, that little will be left to the theory of speech acts" (1983, p. 278). One might call this the dissolution thesis with respect to speech act theory. Jacobs (personal communication) advances a similar view that speech act theory will dissolve into a more fine-grained analysis of rationality, intention and inference.

[14] Learning the inverses of many-to-one functions is a classic and difficult problem in engineering, one for which good solutions are only just being devised (see e.g., Jordan & Rumelhart, 1990).

[15] The founders of the classical plan-based framework are well aware of these problems. In fact, recent work (e.g., Cohen & Levesque, 1990) recommends that the focus be taken away from illocutionary force and instead turned toward patterns of intentions and their logical interrelationships.

indefiniteness of the form/function relation. If linguistic forms can not be unambiguously indexed by functional descriptions, then the proposed functional indexing schemes must inevitably prove unsatisfactory.

To avoid misinterpreting my argument, it is important to note that the failure of functional indexing as a strategy for instantiating abstract message descriptions is not tied directly to weaknesses of speech act theory. As the next section describes, the problem is deeper and is due in part to a fundamental (though not widely appreciated) property of messages, namely their indefinite describability (D. O'Keefe, 1989, 1991).

The Indefinite Describability of Messages

The crux of this argument is that messages are indefinitely describable. There is no one true description of a message (or any linguistic form):

> there is no intrinsically correct system for describing messages - -that is, there is no one intrinsically correct descriptive vocabulary for messages (utterances, texts, etc.), no one correct message category system, no one correct way of distinguishing and classifying messages. There are a great many descriptive vocabularies for communicative conduct, and no one of them is the correct descriptive vocabulary. (D. O'Keefe, 1987, p. 4)

Clearly, appreciating this fact about messages does not immediately damn speech act theory (or any other functional description) as a descriptive vocabulary. There are an indefinite number of descriptions one might attach to utterances, each more or less appropriate to some empirical question or other. The point is that no single descriptive vocabulary can exhaustively describe an utterance: "no category system can embody all of the possible dimensions of differences among interactional behaviors (because no category system can simultaneously orient to all of the possible message features)" (O'Keefe, 1987, p. 13). One must therefore consider the possibility that features other than those captured by speech act theory (or by Hovy's, 1990, rhetorical features, or by any other single vocabulary) might be systematically related to some aspects of the planning process being modelled.

The point is not to say a given description (e.g., speech act theory) is the wrong functional description to apply to messages. The descriptive framework provided by speech act theory has and continues to demonstrate its utility in a variety of research settings. Rather, the point is to suggest that the reification of any one descriptive vocabulary, especially in cognitive theorizing, is counterproductive. O'Keefe (D. O'Keefe, 1987) discusses one pattern of reasoning that illustrates precisely this mistake.

The mistake is to presume that, because some descriptive vocabulary seems to capture empirical regularities in communicative behavior, the terms of the vocabulary must themselves have a causal role to play in the cognitive processes that produce the behavior. As D. O'Keefe notes, "the very ability of a descriptive vocabulary to identify behavioral regularities is commonly taken as evidence that something like the elements of that vocabulary are already represented in the cognitive systems of social actors" (1987, p. 29). This kind of reification is obviously at work in speech act and other pragmatic models of language production. For example, speech act theory seems to capture a lot about the web of social commitments that surrounds everyday actions. Therefore, speech act categories are put in the head to generate the patterns of behavior they explain.

There are problems with this way of thinking. First of all, any category system (or other descriptive vocabulary) will run into the same form/function problems that so plagues speech act theory. For instance, "any message token (any actual message) exemplifies an indefinitely large number of message types" (D. O'Keefe, 1987, p. 11). The equivalent problem with regard to speech act theory is that, in the right context, an utterance can perform an indefinite number of acts.

There are two conclusions to draw from this line of argument. First, the typification of messages is a necessary and valid part of capturing generalizations about messages. Second, the analysts' types should not be reified and given a causal role in message production. When they are, a type-token relationship is introduced that only an untenable functional indexing scheme can solve. The section that follows questions the standard model's reliance on stable associations between types of situations and types of messages, arguing that message variability makes such stable means-ends relations seem quite implausible.

The Typification of Means-Ends Relations

In most models of message production, the problem of deciding what to say is solved in the abstract by a fixed set of rules that associates types of situations with types of messages (e.g., "In a persuasive situation, emphasize advantages to the other," "If the goal is to commit to the truth of p, assert p," etc.). These rules codify the relationship between social ends and linguistic means. The problem with this approach is that means-ends relationships are not uniform. Across individuals and over the life-span of development, the evidence reveals diversity in means-ends relations. Models of message production must explain this diversity. Until now it has largely been ignored.

A great deal of recent research has shown that different people will produce very different messages in response to the same complex communication situation (for reviews see Dillard, 1990; O'Keefe, in press a; O'Keefe & Delia, 1988; Seibold, Cantrill, & Meyers, 1985; Tracy, 1991). Explaining this diversity is a major aim of contemporary communication research. At a minimum, there are two types of diversity one encounters in the study of message production: (a) The same person may produce different messages in response to different situations; and (b) Different people may produce different messages in response to the same situation.[16] The first type of diversity is more frequently studied than the second and is typically explained with reference to higher-order diversity at the level of goals. Different situations involve different sets of goals, and hence lead to the production of different messages. Such an account of diversity is well illustrated in the work of Hovy (1990), for example.

Although its existence has important implications for theories and models of message production, the second type of diversity has until recently been unappreciated by most analysts. O'Keefe (Lambert & O'Keefe, 1990; O'Keefe, 1988, 1990a, 1990b) has argued that this sort of diversity originates not in different goals but in different styles of reasoning ("message design logics") linking goals to verbal means for achieving those goals. Before considering the details of O'Keefe's (1988) theory, the core evidence of message diversity is reviewed.

Functional Message Production

Functional communication skills enable people to use linguistic resources to do the everyday work of social interaction (e.g., comforting, persuading, regulating, informing, etc.). These skills are not evenly distributed across the population of language users; it is self-evident that some people are better than others at composing messages to achieve everyday social goals. According to O'Keefe & Delia (1988, p. 70), "Individual differences in performance have been observed across a range of communicative tasks, involving such diverse processes as reference and description, persuasion and behavior regulation, comforting and interpersonal problem solving, and conflict management and negotiation." Accounting for these systematic individual differences is an important task for communication researchers both at a theoretical and a practical level. Any analysis failing to

[16] There is a third type of diversity not discussed here. The same person may respond to the same situation differently at different times.

offer an explanation of functional diversity in messages would be found wanting.

Examples of Functional Diversity in Message Designs

Developmental studies. Early work in the constructivist research tradition provides clear examples functional diversity in message designs. Delia, Kline, and Burleson (1979), for example, asked children in kindergarten through twelfth grade to produce a variety of persuasive messages. Subjects responded to several hypothetical persuasion situations (ask your mother to have a slumber party, ask a neighbor to take care of a stray puppy, etc.). Not surprisingly, Delia et al. (1979, p. 248-249) found that persuasive appeals increased in sophistication with age. For example, the youngest subjects produced unelaborated requests ("Can I have a party?", "Could you keep this dog?"), or simple statements of personal wants, needs or desires ("I want a party"), or repeated simple appeals ("Please Mommy please, can I?"). Slightly older children produced messages that elaborated the speaker's needs ("I've never had a slumber party, and I really want to have one") as well as messages that elaborated the needs of involved third parties ("This poor puppy needs a place to stay", "My friends would really enjoy a slumber party"). Still more sophisticated kids produced messages that anticipated counter-arguments ("Could you keep this puppy, it is old enough to be outside. It won't make a big mess"). Finally, the oldest and most competent subjects produced messages that included explicit mentions of advantages to the other ("Dogs are great playmates for kids", "You live alone, and you need a good watchdog like this one"). In a developmental context, this pattern of results illustrates how a developing appreciation for the perspective of the other, combined with increasingly more elaborate conceptions of messages and their effects, can be responsible for the production of messages that are adapted to their tasks in quite distinct ways.

Non-developmental studies. It is hardly surprising that children of different ages produce different sorts of messages in response to the same tasks. If this were the only kind of observed message diversity, standard accounts of message-level production could be excused for ignoring it. Such theories, one might argue, are intended only to model mature communicative performance, once competence has levelled off at some steady state. In point of fact, however, functional diversity of message designs is also evident in messages produced by age-homogeneous populations. In a recent study, O'Keefe and Lambert (1989) asked college undergraduates how they would respond to an old friend, Terry, who had just broken a date to go to a movie. Some of the messages produced in response to this task are reproduced as (9)-(13) below:

(9) Terry, you're a jerk! You're not a real friend! I don't want to be your friend anymore if you're never going to treat me like a friend. You're insensitive and uncaring! Goodbye!

(10) Sorry, I don't think we should plan on going tomorrow night. You've been backing out on everything lately, so why don't you give me a call sometime when you are free and we'll stop planning things. I suggest you start to think of others and not just yourself a little bit more.

(11) I understand you've got a lot to do, but if you can't fit our friendship into your schedule, I wish you'd just say so.

(12) Yeah, I guess. But I'd like you to know that I'm really getting tired of you backing out of our plans. I mean, why make plans when you cancel all of them? I'd say just call me when you find yourself with nothing to do, except then I'd never speak to you.

(13) Terry, are you too busy for me or what? You're a great friend and I really enjoy being with you. That's why I keep planning things for us to do. You either are not admitting to me that you do not want to spend time with me or you're kind of irresponsible in terms of scheduling. Which is it?

(14) I feel hurt that you've cancelled all of our previous activities. I should tell you "no I won't see the movie tomorrow," but you're still my friend and I do enjoy spending time with you. So why don't you give me a call tomorrow night. If we're both free, we'll see the movie. This way, one of us won't be waiting for the other if something comes up.

Notice first of all that a complex situation like this one (the Broken Date situation) elicits substantial variation in message contents. Message (9), for instance, insults Terry, scolds him/her for previous bad behavior, and dismisses Terry as a friend. Messages (11) and (12), while still voicing frustration with Terry, abstain from explicit insults and instead attempt to change Terry's behavior by reminding Terry about the norms governing friendship and the scheduling of dates. Messages (13) and (14), by contrast, remind Terry of his/her value as a friend. Message (14) also differs from the others in its focus on the future, exploring alternative scheduling arrangements to suit both parties' needs.

The logic of message design. These messages illustrate that individual differences in message designs exist even in age-homogeneous populations, and they persist well into adulthood. Different people respond to the "same" situation in dramatically different ways. O'Keefe (1988) argues that the different responses are symptomatic of fundamental differences in the way people reason from goals to verbal means for achieving those goals. These different patterns of reasoning from goals to verbal means O'Keefe (1988) terms "message design logics." Message design logics are ontogenetically plastic cognitive structures that link goals to message contents. They implicitly embody message producers' beliefs and theories about the nature of language, social relations, messages, and message effects. O'Keefe (1988) argues that each message producer relies on one of (at least) three developmentally ordered, internally consistent, and readily identifiable message design logics. The three design logics, in order of their appearance in development, are called expressive, conventional, and rhetorical (see O'Keefe, 1988).

An expressive message design logic is built on the presumption that language is primarily a means for expressing the speaker's own thoughts and emotions (O'Keefe, 1988). Within the expressive logic, messages are seen merely as vehicles for expression. Speakers generating messages out of an expressive logic may know that messages have effects on hearers, but messages are not designed to achieve any specific effect. Rather than instrumentality, openness, honesty, and unimpeded clarity of expression are the most salient evaluative dimensions for messages designed within an expressive logic. This lack of focus on the instrumental potential of messages leads expressive message producers to attend exclusively to the literal and direct content of their own and others messages (O'Keefe, 1988). Because the principles of coherence organizing expressive messages are idiosyncratic, they may take on a semi-coherent, stream of consciousness quality, repeating ideas and jumping from idea to idea seemingly at random. The primary mechanisms for adaptation within the expressive logic are editing and deception, reflecting a belief that one can either not say what one is thinking or one can lie about what one is feeling and thinking.

A conventional message design logic reflects the belief that language is a kind of game played cooperatively by socially defined rules (O'Keefe, 1988). In contrast to the purely phatic function of expressive message designs, conventional messages are straightforwardly instrumental and intentionally designed to achieve specific effects. Principles linking messages and message effects, such as Grice's (1975) maxims of cooperation and Searle's (1969) felicity conditions, are conventionally specified. Within a conventional message design logic, message elements are selected for expression because expressing those particular elements in the given context counts as performing a particular act, and the performance of a particular act is conventionally associated with certain effects (O'Keefe,

1988). Thus, the effect of a message in context is mediated by conventional principles (e.g., felicity conditions). Similar conventional principles (e.g., adjacency pair organization) govern the coherence of conventional messages. Within the conventional logic, social context, including the structure of rights and obligations attached to status-bound social relations, is taken as given. The primary evaluative dimension operating within the conventional logic pertains to the appropriateness of a message given the context (Lambert & O'Keefe, 1990; O'Keefe, 1988). The primary means for functional adaptation within the conventional logic are indirectness and redress (Brown & Levinson, 1987). The vast majority of plan-based models reviewed here, especially those founded on speech act theory (Appelt, 1986; Cohen & Perrault, 1979) model only the conventional logic of message design.[17]

A rhetorical message design logic "reflects a view of communication as the creation and negotiation of social selves and situations" (Lambert & O'Keefe, 1990, p. 6). Rather than treating context as fixed, given, and as determining what is appropriate to express, the rhetorical logic sees roles and contexts as the products of language use (for entire theories founded on this view, see, e.g., Blumer, 1969; Goffman, 1967). Whereas the conventional logic links messages in context to their conventionally specified effects, a rhetorical logic captures "the ways in which communicative choice and language style convey character, attitude, and definition of the situation" (O'Keefe, 1988, p. 87). Rhetorical messages are coherent to the extent that they portray a consistent and consensually agreed upon definition of roles, character-identities and events. Whereas conventional messages are coherent when they count as legal moves in a given language game, rhetorical messages must reflect a sensible narrative structure to be coherent.

Message producers operating within a rhetorical logic intend for their messages to be interpreted as signals of character, motivation, etc., and they rely on the same kind of deep interpretive strategy in understanding others' messages. Rhetorical reasoning has not been adequately addressed by computer models of message production, and there is reason to suspect that modelling rhetorical patterns of reasoning within the framework of standard planning models may not be a very promising project. Before analyzing the difficulties one might encounter in modelling the three design logics, evidence is reviewed that demonstrates the differential effectiveness of messages generated out of the three different logics.

Consequences of functional diversity. Some messages work better than others. That is, messages can be adapted to tasks in a variety of ways, and not all adaptations are equally effective. For example, Burleson and others (Applegate & Delia, 1980; Burleson, 1984; Burleson & Samter, 1985) found that some comforting messages were much more effective than others. Burleson (1984) asked subjects what they would say to comfort another child who had not been invited to a friend's party. Messages that challenged the legitimacy of the target's feelings ("There's nothing to be upset about") were rated as less sensitive than messages that recognized the target's perspective ("I'm sorry you did not get invited"). Messages that explicitly recognized and elaborated the target's perspective (I know you feel bad, but you're my friend and lots of people like you") were rated as most sensitive. Investigating the perceived effectiveness of messages generated from different design logics, O'Keefe and McCornack (1987) found that rhetorical messages were perceived as more polite and more effective than conventional or expressive messages, and that raters were more likely to want to spend time with rhetorical message producers than with either

[17] I mean that such models capture nothing more than conventional reasoning. In fact, in their failure to address indirectness (see e.g., Appelt, 1986), such models fall far short of full conventional functioning.

conventional or expressive message producers.

The literature on message effectiveness shows that functional diversity has important practical consequences. In view of its theoretical and practical significance, models of message production must be held accountable to the evidence of message diversity. I contend in the next section that current plan-based models fall far short when held accountable to these data.

Typification and Message Diversity

The existence of non-developmental individual differences strongly suggests that people do not share a common mental model of message effects. Thus any model that internalizes a single functional description of messages and statically associates a situation-type with each message-type will fail to reproduce the variability that is a central theoretical concern for communication researchers. On these grounds the reification of fixed abstract characterizations of messages and the explicit indexing of acts by effects must be rejected.

When diversity is explained on a design logic account, this point becomes especially clear: forms only have functions within a given design logic. If message variability is a consequence of different people exploiting fundamentally different means-ends relations, then it is inappropriate to privilege (reify, mechanize, etc.) a fixed set of means-ends relations. Rather, a cognitive theory of communication should seek to discover what kind of structures and processes would naturally produce the sort of means-ends diversity observed in compliance-gaining research and elsewhere. At present, it is not clear either how a plan-based model could capture the different styles of reasoning that are captured by the three design logics.

Constructive operators and social structure. Most models of message-level language production are based on the metaphor of the language game. On this metaphor, language is governed by rules that are conventionally agreed upon rather than necessary or natural. So if communication is like a chess game, utterances (conceived as acts) are the moves. The board itself is the social world, and conventional principles (e.g., felicity conditions) specify the legal moves in the game and their consequences. This is a superb rendering of a conventional design logic, but how does expressive or rhetorical reasoning fit into this image?

The expressive communicator is not even aware that a game is being played. S/he arranges pieces on the board for what seem like purely artistic purposes, without regard to the potential instrumentality of his/her actions under some interpretation (e.g., the game interpretation). The conventional communicator moves about the board according to the rules and in pursuit of standard objectives (e.g., checkmate). S/he may move the pieces strategically according to some more or less rational plan for winning the game. In contrast, the rhetorical communicator feels free to redraw the grid to her own specifications, or to rewrite the rules of the game, sometimes moving pieces in strange and unconventional ways, sometimes claiming to be engaged in a different game altogether (e.g., checkers). This is all to say that rhetoricals and expressives use operators (they move about in social, pragmatic, discursive space) in unfamiliar ways, ways that conventional planners have yet to fully explore.

Granted that there may be a whole host of operators as yet undescribed, there are a couple of ways to proceed. The first and most obvious would be to get on with the work of locating and defining the novel operators that expressives and rhetoricals use. There is reason to suspect that such a project, conceived in traditional terms, is unlikely to succeed. The very idea of an operator, a move through problem space with fixed preconditions and consequences, is tied to the conventional root metaphor for communicative action in the social world.

The conventional metaphor does not capture the kind of problem space in which the expressive operates. Expressive messages are planned not as actions in the service of goals, but as expressions for the sake of expression. Whereas the primary telos in conventional planning is taken to be instrumentality of utterance-as-action, the expressive telos is for clarity and accuracy of representation in expressions. It is not clear what it would mean to plan expressions with the instrumental objective of accurately representing thought in words. What would be the subgoals in such a plan?

The differences between design logics may be glossed in terms of optimizing with respect to different objectives. The expressive logic generates messages that are planned to be optimal with respect to accuracy of representation of thoughts in expressions. Messages generated out of an conventional logic are planned to be optimal with respect to (conventional) instrumental goals. Finally, the rhetorical design principles produce messages that are optimal with respect to the degree of consensus they achieve.

But the rhetorical logic also seems to differ from the image embodied in the conventional metaphor, in that the former makes use of what might be called constructive operators. Continuing the game-board metaphor from above, notice that rhetoricals do not respect the rules and boundaries of the game being played. The conventional reasons only within the narrowly circumscribed possibilities defined in advance by the rules of the game being played. In the context of a chess match, the conventional sees a pawn as a pawn and moves it accordingly, but the rhetorical is not similarly constrained. Rhetoricals may use language to create new symbolic spaces that serve as shortcuts to their goals. The rhetorical will use the pawn as a queen, or better still, shift the game to checkers, if such a redefinition serves the current rhetorical purpose.

Given this kind of dynamic, it is reasonable to wonder whether constructive operators (e.g., the "explicit contextualizing phrases" of O'Keefe, 1988) can be defined in terms of preconditions and consequences as would be necessary in a plan-based scheme. The notion of meeting preconditions takes for granted a fixed and stable set of conditions at any given moment. It is precisely these conditions (the definition of the situation) that constructive operators change, redefine, recreate, etc. Such a rule would have to look something like the following:

If conditions $C_1, C_2, \ldots, C_n$ obtain,

Then apply constructive operator O_1 to create $C_1', C_2', \ldots, C_n'$.

But it is not clear how one could follow such a rule. If it is sometimes okay to redefine the current situation, how does one know when to stop recontextualizing and respond the current set of conditions in and of themselves? This is the sort of dilemma that arises in trying to accommodate the facts of message diversity within a plan-based framework.

Summary

The critique of plan-based approaches depends on two related lines of argument. First, plan-based approaches exhibit an overreliance on a kind of stable form/function distinction that can not be supported empirically. Using speech act theory as an example (Levinson, 1979a, 1983), it was shown that the form/function distinction breaks down in even the most mundane and trivial circumstances. The existence and pervasiveness of indirectness frustrates any neat form/function mapping. Utterance units can not be determined independently of functional classification. The same utterance unit may instantiate many acts simultaneously, and no explicit procedure exists for assigning utterance units to act categories (Levinson, 1979). Concerns about the indefinite describability of messages further undermine any lingering hopes about discovering a single functional description of messages or message features (D. O'Keefe, 1987). Together, these observations were used to argue against the possibility of constructing a functional indexing scheme that could be used to map

from acts to utterances.

The second main line of argument against traditional plan-based models of message production asserts that populations of messages exhibit functional diversity in their design, and that existing models do not address this dimension of diversity. Functional diversity in message designs was illustrated with a variety of examples, and was shown to have important practical consequences in terms of message effectiveness.

Traditional models seem ill equipped to deal with this kind of diversity for three reasons: (a) they rely on typifications of situations (goals); (b) they rely on typifications of messages; and (c) they rely on stable associations between situation-types and message-types. Such systems, it was argued, do not account for indirectness, the multifunctionality of messages, or diversity in message designs. On these grounds, plan-based models must be rejected as inadequate explanatory models of message production.

The remainder of this thesis attempts to construct a model of message production that does not fall victim to these criticisms. Chapter 4 provides an overview of the model. Chapter 5 describes a strategy for avoiding the typification of situations. Chapter 6 describes a strategy for avoiding the typification of messages, and chapter 7 describes a strategy for avoiding the typification of means-ends relations.

CHAPTER FOUR
LOCAL MANAGEMENT AND MESSAGE DESIGN: A PDP VIEW

Although rational goal analyses of the sort offered by O'Keefe (1988; see also Brown & Levinson, 1987) are not themselves theories of message production, they lend themselves to implementation in models of the message production process. One chief objective of the current project is to develop such an implemented model of the message production process based on O'Keefe's (1988) theory.

There are several reasons why it would be desirable to have such a model. First, and perhaps most important, many of the claims O'Keefe has made about message design logics can be expressed and tested through an implemented model of message production processes. For example, she has claimed that differences in design logic result from the progressive acquisition and integration of knowledge about the mapping between communicative forms and social structures. One way to develop and test this claim is to simulate learning about some communicative task and observe the performance of the learning system as it absorbs information about the task.

Second, development of detailed models of message production should provide information useful to the design of training and intervention programs. Differences in skill at message design are consequential for individuals, so it is important to determine what kinds of performance individuals are capable of (or can be made capable of). Development of an implemented model would directly facilitate the design of training procedures.

Third, development of such models provides the best current method for cross-validating message analysis systems based on message design logic. A variety of systems have been developed for classifying messages in terms of the design logic they seem to be based on. Until recently such systems could only be justified in terms of their face validity. However, recent work by O'Keefe and Lambert (1989) has shown that the articulation of detailed message production models can promote development of procedures for assessing differences in communication-relevant beliefs, which in turn can be used as criteria for assessing the validity of message analysis systems.

Finally, most models of message production are based on the assumption that communicative rationality is universal, both in the sense of being held in common by all communicators and in the sense of being constituted independently of particular communication situations. O'Keefe's (1988) theory is anti-universalist in both respects. Developing an implemented model based on this theory would provide a useful demonstration of the fact that universalism is not a necessary property of an implementable rational analysis.

In this chapter I (a) argue for a situation-based representation of message organization; (b) argue for a distributed representation of communicator goals; and (c) present a detailed model of message design.

A Situation-Based Representation of Message Structure

The most common approach to representing messages, an approach that cuts across both observational studies of message diversity and efforts to implement models of message production, involves characterizing message abstractly as types of action (e.g., speech acts) or types of message strategies (e.g., Brown & Levinson, 1987; for a more detailed discussion of the concept of "strategy," see O'Keefe, 1989). The classification of messages into act- or strategy-types necessarily involves the abstraction of the function of a message from the particular thoughts that the message expresses.

This abstractive move--away from the embodied content of the message and toward an

analysis of the message function--is crucial to a theory of message adaptation and effects. Although it may ultimately prove impossible to have a general theory of message design, in which types of message designs are related to types of outcomes, the effort to construct such general theories rests on the assumption that messages have effects as types. Consequently, one cannot avoid the development of rational analyses of message form: such an analysis is constitutive of a theory of adaptation.

However, the fact that a typification of messages is helpful for understanding the rational properties of their design does not imply that the categories of typification literally play a role in generating the message. Consider, for example, the variable "message length." Imagine that we find that messages differ in length and that longer messages are perceived as more effective than shorter ones. This would hardly be taken as evidence that there is a specific procedure in people's minds that is responsible for generating the feature "length." The point is that typifications of messages, which are indispensable from the standpoint of understanding message functions, may be irrelevant from the standpoint of modelling the processes that generate messages.

In fact, O'Keefe's (1988) argument presumes that no one functional description of messages can possibly be adequate as the basis for a model of message production across individuals; this presumption is derivable from the core assertion that each message design logic reflects a different way of conceptualizing the job a message does and what the message's adaptive properties are. Since conceptions of communication function are particular to design logics, any given functional description will reflect the application of some one logic. A different logic would generate a different, and equally cogent, functional account. Hence, the project of developing an implementable model of message production based on the analysis of message design logic requires finding a different way of representing message structure that does not presuppose, a priori, a functional vocabulary.

The strategy adopted here for developing such a representation involves characterizing messages as a string of expressed thoughts (for similar strategies, see Chafe, 1979; Levy, 1979). This kind of representation is necessarily situation-specific, since very similar-looking specific thought units may perform very different functional roles in different situations (e.g., "I like your sweater" might function as a compliment in one situation and a request in another) and very different-looking thought units may play the same functional role across a range of situations (e.g., what is needed to give social support to another may differ considerably from one situation to the next).

So rather than attempting to develop an abstract representation of message structure, this chapter attempts to develop a representation that preserves the situated character of individuals' expressed thoughts. The details of the methodological application of this approach are discussed below.

Distributing the Representation of Communicator Goals

Just as I have been led to represent messages as strings of situation-specific types of thoughts, so I have been led to adopt a distributed representation of communicator goals. Most current treatments of the message production process conceptualize goals as symbolic atoms with no internal structure of their own (for examples of this strategy see the chapters in Cohen et al., 1990, or Hovy, 1990; for a general discussion and critique of this strategy, see Lave, 1988; Suchman, 1987). But a view of goals as specifiable atoms that motivate and organize message planning is at base inconsistent with the premises of O'Keefe's (1988) theory.

A message design logic embodies both a conception of what goals communication can serve and a conception of how it serves those goals. Hence, there is no general set of functions messages can serve apart from design logic. The assignation of functions to

messages can only be done within a particular design logic. Since the meaning of message forms is distributed throughout an individual's knowledge of communication, and since message design logic is embodied in the network of form-meaning relationships in an individual's system, and since functions are defined by the individual within his or her own design logic, the proposed model relies on a distributed representation of goals.

Hence, the ground of communication, its goals, are conceptualized as distributed in the pattern of thoughts about a situation that an individual has as he or she produces a message. Again, the methodological application of this principle is discussed below.

A Model of Message Production

This section describes an initial attempt to specify, within the constraints imposed by the theory of message design logic, the cognitive structures and processes implicated in message production. The model is based on 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 the 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 one's available knowledge about a topic is organized for expression during a given turn at talk.

Herrmann's Basic Model of Speech Production

Herrmann's (1983) model is a multi-step, cognitive model of language production wherein "production is conceptualized as the transformation of nonverbal cognitive content into verbal utterances" (Hoppe-Graff et al., 1985, p. 82). The first step in the model calls for the construction of a situation representation by the situation-interpreting program (see Figure 3). The resulting situation representation contains both information about relevant situation parameters as well as procedural and declarative knowledge retrieved by the situation interpreting program. Situation parameters refer to those features of the immediate context that will impact on 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 (Schank & Abelson, 1977) 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.

The propositional base. 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. Thus "the propositional bases of the utterances activated by speakers (or the 'meaning meant') represent declarative knowledge; its activation and implementation on behalf of the adopted goal may require the use of procedural knowledge" (Herrmann, 1983, p. 36).

The semantic input. 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, p. 84), "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." 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). Lower level processes (articulation, linearization, etc.) that produce the overt utterance are not discussed.

Summary of the model. 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. The project adopted here is to model the second of these steps, in which the propositional base is transformed by selection into the semantic input. In the next section it is argued that this selection is accomplished by relevance selectors, associative networks that embody the situation specific functional mapping form propositional base to semantic input.

Selection of Semantic Input

It is neither possible nor necessary for a message producer to utter every proposition in his or her currently active propositional base. Therefore, a mechanism is needed to select the subset of propositions that are relevant for expression. Conceived as such, the basic principle of all message production is simply "say what's relevant" (see Sperber & Wilson, 1986, for a version of this thesis). On this view, message variation is due to the operation of different senses of relevance. It is argued here that the knowledge structures that embody one's sense of relevance are learned from concrete experience in many situations. People learn from experience in concrete situations what clusters of propositions in their propositional base count as relevant propositions to express. From this brief description, it should be obvious that the existence of senses of relevance presumes the existence of a powerful learning mechanism. The next section describes a candidate mechanism to account for the acquisition and use of senses of relevance.

PDP and associative learning. It is hypothesized that a general associative learning mechanism can be used to model the acquisition and to simulate the performance of a situation-specific sense of relevance. The details of the learning mechanism are based on a class of models described by Rumelhart, Hinton, & McClelland (1986) as supervised pattern associators. A pattern associator is one of a much larger class of so-called connectionist (Feldman & Ballard, 1985) or parallel distributed processing (Rumelhart & McClelland, 1986) models of learning and cognition. Connectionist learning models consist of semantically defined nodes connected by modifiable connection weights. In a pattern associator, the nodes are arranged in input and output layers. The input layer receives information directly from the environment. The output layer receives information which has been mediated by the connection weights. Associated with each node is a level of activation which represents the degree of presence in the current environment of the semantically

defined feature represented by that node.

What makes PDP models interesting, of course, is their ability to learn and generalize from a finite number of examples. When properly trained, such models can learn arbitrary input-output mappings. The models learn by the progressive modification of the connection weights in response to experience (where "experience" often entails feedback from a teacher). Thus, by presenting the system with correct input-output pairs, there exist learning rules which will modify the weights such that a pattern of activation at the input layer brings about the desired pattern of activation at the output layer.

The weight modification procedures typically consist of three steps: the feed-forward step, the error calculation, and the back-propagation step. In the feed forward step, a vector of activation values (from a training example) is clamped on the input layer and multiplied by the (initially random) weight matrix to yield values at the output layer. The resulting output is compared with the desired (teaching) output corresponding to that input pattern and the difference is calculated as an error signal. The third step involves propagating the error signal back through the network and modifying the connection weights accordingly (the specific rules are described in greater detail in the methods sections below).

In the mapping modelled here, the level of activation at the input layer represents the degree of presence of a single proposition in the currently active propositional base. The level of activation of each node at the output layer represents the degree of presence of a single message element in the currently active semantic input.

Summary

One important task for communication theory is to explain diversity in the functional design of messages. Research in this tradition has sought to explain three central facts: (a) that people produce different messages in response to the same situation; (b) that these messages differ in their effectiveness (broadly construed); and (c) that competence is acquired developmentally. Within this tradition, O'Keefe (1988) has developed an image of message production as crucially involving, in every instance, a situation-specific sense of relevance. On this account, the development of functional competence is seen as the cumulative result of concrete experience in specific situations. Such experience enables social actors to learn the mapping from what they are thinking to what they should say. Metaphorically, this experience is captured as s function that maps the propositional base onto the semantic input (thoughts to messages).

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 features of the situation as relevant for expression?" In this view goals are patterns of thoughts about the situation. It is a fundamental premise of the view being proposed that there is no straightforward mapping from abstract goals or strategies to message content (O'Keefe, 1990b).

The next three chapters begin to build an empirical case for this image of message production. Chapter 5 describes a protocol-analysis investigation of people's thoughts as they respond to hypothetical communication tasks. Chapter 6 describes a system for representing messages as activation vectors. Chapter 7 describes a PDP model of the mapping of propositional base to semantic input (i.e., thoughts to messages).

CHAPTER FIVE
DEVELOPMENT AND VALIDATION OF THE THOUGHT CHECKLIST

To implement the model of message production outlined in the previous chapter, it was necessary to develop a method for thought elicitation that would allow patterns of reported thoughts to be represented as numeric vectors. Toward this end, O'Keefe and Lambert (1989) developed a thought checklist task to accompany the Broken Date hypothetical communication scenario. The thought checklist consists of 122 thoughts, separated into different conceptual categories (e.g., thoughts about own feelings, thoughts about other's feelings, etc.).[18] The Checklist task requires each subject to put an "X" next to the thoughts that s/he had as s/he read and responded to the the Broken Date Problem. The thoughts on the Checklist were compiled from a pilot study reported in O'Keefe and Lambert (1989). The thought checklist was used successfully in O'Keefe and Lambert's (1989) study of message design logic-related differences in patterns of reported thoughts, but questions remain about the validity of the thought checklist as a measure of thoughts: (a) Does the checklist task lead people to report more or fewer thoughts than they would in a concurrent or on-line elicitation technique? (b) Does the checklist itself provide people with substantively different ideas than they would have had on their own? (c) Is the checklist a reasonably comprehensive list of thoughts about the Broken Date Problem?

The purpose of this chapter is to cross-validate the thought checklist as a method for assessing people's thoughts as they read and respond to hypothetical communication scenarios. The cross-validation is accomplished through a detailed comparison of thought checklist data and data collected in think-aloud and retrospective protocols. Protocol analysis is a well-established technique for measuring the contents of cognition during problem solving (Ericsson & Simon, 1984; Hayes & Flower, 1980). Substantial similarity between the two methods in the number and type of thoughts reported will be taken as evidence of the thought checklist's validity. The lack of such overlap will be taken as evidence of a need to modify or revise the thought checklist.

Methods and Design

The aim of this research is to understand the relationship between situationally relevant thoughts and produced messages. In this investigation, as in prior work on the logic of message design (O'Keefe, 1988; O'Keefe & Lambert, 1989), messages were elicited from subjects using a hypothetical written communication task. This study compared two different methods of eliciting message producers' situationally relevant thoughts: thought checklist or think-aloud protocol. Differences in the number and type of thoughts reported using these two methods were examined.

Subjects

The participants in this study were $N = 145$ students at a large midwestern university. The subjects were divided into two conditions: Thought checklist data were collected from 114 participants. Incomplete data from one subject were thrown out, resulting in $n = 113$ for the thought checklist condition. Thirty-seven students participated in the protocol condition. Incomplete data from 5 subjects were excluded, leaving $n = 32$ subjects for the protocol condition. All participants were enrolled in an introductory communication course and received course credit for their participation.

[18] An example of the instrument used in this study is reproduced as Appendix A. Due to a typographical error, one thought ("Terry is overcommitted") appeared twice on the Checklist. Responses to the second occurrence of this thought have been omitted in all subsequent analyses, leaving a total of 121 thoughts on the thought checklist.

Thought Checklist Condition

Participants assigned to the thought checklist condition completed a written questionnaire in which they first responded to the Group Leader Problem and then to the Broken Date Problem.

Group leader problem. All participants in the thought checklist condition were first asked to respond to the Group Leader Problem (see O'Keefe, 1988), in which they were asked what they would say to a fellow group member (Ron) who needs more time to finish a project. Participants were asked to write exactly what they would say rather than describing the the general action they would take.

Broken date problem. After writing their responses to the Group Leader Problem, participants were asked to read and respond to the Broken Date Problem (see O'Keefe & Lambert, 1989; Newell & Stutman, 1982). In this problem, participants are asked what they would say to a friend (Terry) who calls to reschedule a date after recently cancelling several previous dates. Again, subjects were told to write exactly what they would say, rather than describing the general action they would take.

After writing a message for the Broken Date Problem, participants were asked to complete the thought checklist task for the the Broken Date Problem. Participants were asked to go through the list of 121 thoughts, putting an "X" next to those thoughts that occurred to them as they read and responded to the Broken Date situation. Participants were cautioned against checking thoughts that seemed true but that did not occur to them as they responded to the situation.

Summary: Thought checklist condition. As a result of the application of these procedures, each subject in the thought checklist condition produced a Group Leader message, a Broken Date message, and a 121 bit vector of thoughts (1 for each thought selected, 0 for all other thoughts).

Protocol Condition

Participants assigned to the protocol condition completed a think-aloud protocol and then a retrospective protocol while reading and responding to each of the two problems (Group Leader and Broken Date). Descriptions of the scenarios were printed in multi-page booklets, with one-sentence per page. The Group Leader booklet was 15 pages long; the Broken Date booklet was nine pages. By assigning one sentence to each page, it was possible to identify the juncture in the scenario story at which thoughts occurred.

Participants were asked to read through their booklets and to speak aloud any thoughts they were having as they read the booklet. The instructions for the think-aloud portion of the protocol task were as follows:

> In this experiment we are interested in learning what people think about when they respond to interpersonal communication situations. In order to do this, I am going to ask you to THINK ALOUD as you read about an interpersonal communication situation and respond to it. What I mean by think aloud is that I want you to tell me EVERYTHING you are thinking from the time you first begin reading about the situation until you have finished responding to the situation. Don't try to plan out what you say or try to explain to me what you are saying. Just act as if you are alone in the room speaking to yourself. It is most important that you KEEP TALKING. If you are silent for any long period of time I will remind you to keep talking. Do you understand what I want you to do?
>
> Good, now we will begin with a practice situation. This first situation is just to let you become more comfortable thinking aloud. Please read each sentence of the

description aloud so I know where you are as you tell me your thoughts.[19]

The last page of each booklet asked participants to write the message they would produce in response to the scenario. As in the thought checklist condition, participants were asked to write exactly what they would say, rather than give a description of the general action they would take. Participants were instructed to continue thinking aloud as they produced the message, so that the experimenter could record any thoughts that occurred to them as they wrote their messages. Using these procedures, participants produced think-aloud protocols during the reading of the scenario and the writing of the message for each of the two problems.

After a participant had finished responding to a problem, the booklet for that problem was collected and the participant was asked to provide a retrospective protocol. Participants were told:

> Good, now I want to see how much you can remember about what you were thinking from the time you read about the situation until the time you completed your response. We are interested in what you can actually remember rather than what you think you must have thought. If possible, I would like you to tell me about your memories in the sequence in which they occurred while you were reading and responding to the situation. Please tell me if you are uncertain about any of your memories. I don't want you to think about responding to the situation again, just report all that you remember thinking about as you read and responded to the situation. Now, tell me what you remember.

Hence, for each problem, participants generated (a) a message, (b) a think-aloud protocol for the scenario and message, and (c) a retrospective protocol for the scenario and message.

Using these procedures, participants first responded to the Group Leader Problem and then to the Broken Date Problem. The Group Leader Problem was used as a training task: the experimenter monitored the response to insure that the participant understood and executed the think-aloud procedure. After completing the Group Leader problem and before beginning the Broken Date Problem, participants were asked to discuss the classes they were currently enrolled in. This distractor task took approximately five minutes and was designed to eliminate any priming effects of the first problem. In all, completion of the two tasks and the distractor task typically took about 30 minutes. All protocols were recorded with a portable cassette recorder through a microphone placed on the desk in front of the participant. Taped protocols were subsequently transcribed.

Analysis and Results

Quantity of Reported Thoughts

The first question this study was designed to address was whether the thought checklist task led people to report more or fewer thoughts than they would in a concurrent or on-line elicitation technique? To answer this question, the total number of thoughts reported by each participant was tallied. For the Thought Checklist condition, the total number of thoughts selected by each participant was tallied. For the Protocol condition, the Broken Date portions of each protocol were first segmented into t-units (independent clauses) according to the procedure outlined in Cooper and Odell (1978). To establish the reliability of the unitizing procedure, two coders independently coded the think-aloud and retrospective protocols of two subjects. Guetzkow's (1950) U, a reliability statistic for unitization, was .036. Values of U less than .10 are generally regarded as acceptable. Separate tallies were then made of the number of thoughts in each subject's think-aloud and retrospective protocols. Thus each subject in the Protocol condition received two scores: a score for number of t-units reported during think-aloud protocol, and a score for number of t-units reported during retrospective

[19] Instructions for the elicitation of protocols are adapted from Ericsson and Simon (1984).

protocol.

The main analysis consisted of a t-test for independent samples of unequal size. The test compared the mean number of thoughts reported in two methods of thought elicitation, think-aloud protocol and thought checklist.[20] There was no significant difference between methods of elicitation $t(143) = 1.66$, $r = .138$, n.s. In a secondary analysis, the mean number of t-units reported during retrospective reports was found to be significantly less than the mean number of t-units reported in talk-aloud protocols $t(31) = 4.88$, $p < .0001$, $r = .378$, and less than the number of t-units in thought checklist reports $t(143) = 5.148$, $p < .001$, $r = .395$. Simple forgetting may explain this difference.

Representativeness of the Thought Checklist

The second question this study was designed to answer was whether the thought checklist itself provides people with substantively different ideas than they would have had on their own? To answer this question, all of the Broken Date protocols were segmented into t-units according to the procedure described above. Each of the resulting 1685 t-units was classified as either present or absent on the current version of the thought checklist, according to a system developed for this investigation.[21] Since the thought checklist consists of standard forms of expression devised by the experimenter, literal matches between protocol-elicited thoughts and items on the Checklist were not expected to predominate. Thus, criteria had to be established for assessing similarity between protocol-elicited thoughts and items on the Checklist. These criteria needed to be flexible enough to capture genuinely synonymous thoughts but rigid enough to be valid and reliable. Early versions of this coding system had flexible criteria for assessing similarity but were difficult to reliably apply. In the end, a fairly rigid set of guidelines was established to increase the system's reliability. To establish the reliability of this coding system, two coders independently coded 43 t-units (the talk-aloud and retrospective protocols of one subject), achieving exact agreement in 81% of the cases; Cohen's (1960) reliability statistic, kappa, was .63. This figure is on the low side of what Cohen deems acceptable. However, one must keep in mind the difficulty of comparing each of more than 1600 tokens with 121 potentially synonymous thoughts.

Of the 1685 t-units in the protocol data, 824 (49%) were classified as present on the current version of the thought checklist. While it is comforting to note that about half of the thoughts reported in talk-aloud and retrospective protocols were already on the thought checklist, still about half (861) of the thoughts reported in protocols were not. These thoughts were subjected to further analysis.

Comprehensiveness of the Thought Checklist

The third question this study was designed to answer was whether the Checklist is a reasonably comprehensive list of thoughts about the Broken Date Problem? That is, are there many thoughts reported in protocols that are not present on the thought checklist? In examining the thoughts coded as absent from the Checklist, it was obvious that a significant proportion were not focused on the Broken Date situation but rather on the mechanics and social context of the protocol task itself. For example, many participants referred to previous experiences that resembled events in the Broken Date Problem (e.g., "I just met a guy named Terry on the tennis team," "This reminds me of my roommate," etc.). Many participants asked questions of the experimenter (e.g., "Is Terry a boy or a girl," "Do I keep reading," etc.). Participants often gave meta-level descriptions of their own thought processes (e.g.,

[20] Number of t-units reported during retrospective protocols was not used included in the ANOVA because it was a within-subjects factor. Separate tests comparing number of t-units reported during retrospective protocol to number of t-units reported in thought checklist and think-aloud protocol are reported below.

[21] The Present/Absent coding system is reproduced as Appendix B.

"That's all I can remember," "I can remember what I thought but not what the situation was," etc.). Generalized subjective reactions were also common (e.g., "That's bad," "That's nice," etc.). In recognition of this pattern of responses, the t-units classified as absent from the current thought checklist were further classified as reflecting a focus either on the Broken Date situation or on the Protocol task itself.[22] To establish the reliability of the coding system developed for this purpose, two coders independently coded 41 t-units and achieved exact agreement in 97% of the cases; Cohen's (1960) kappa was .95.

Protocol-specific thoughts. Of the 861 t-units classified as absent from the thought checklist, 204 (24%) were classified as protocol-specific; the remainder (657 or 76%) were classified as situation-specific. Protocol-specific thoughts should not be expected to appear as options on the thought checklist. If these protocol-specific thoughts are subtracted from the original total of 1685 thoughts, the corrected proportion of protocol-elicited thoughts already on the thought checklist increases to 56% (824 of 1421), with the 95% confidence interval ranging from 37% to 75% (Glass & Hopkins, 1984, p. 280).

Situation-specific thoughts. It still remains to account for the 657 thoughts that were neither on the thought checklist nor specific to the protocol task itself. These thoughts are legitimately missing from the checklist. These thoughts were content analyzed in order to decide which, if any, should be added to subsequent versions of the thought checklist. Before considering specific thoughts, it is important to point out four general findings from the content analysis of thoughts reported in the protocols: (a) missing thoughts fit easily into the existing categories from the thought checklist (e.g., thoughts about own feelings, thoughts about the relationship, thoughts about Terry's motives, etc., see Appendix A), hence the thought checklist did not miss any broad topic area; (b) many missing thoughts might justifiably have been classified as present under a more flexible definition of synonymy (but see the discussion of reliability above); (c) some missing thoughts arose out of idiosyncratic or metaphorical interpretations of the Broken Date situation that could not reasonably have been anticipated on the checklist; and (d) many positively-valenced thoughts about Terry were elicited during the talk-aloud protocol task before information about Terry's bad behavior was introduced; since the thought checklist is taken after the bad behavior is mentioned, positive assessments would presumably be rejected by participants in the thought checklist condition before responding to the checklist.

In spite of these qualifications, the protocol data did suggest some minor revisions to the thought checklist. Perhaps the most deficient category on the thought checklist is the one pertaining to thoughts about the relationship. Several thoughts belonging to this category appeared in the protocols but were not on the Checklist. In fact, the data suggest that we ought to differentiate between thoughts about this friendship and thoughts about friendship in general. For example, several thoughts made reference to the sorry state of the specific friendship with Terry (e.g., "The friendship is in trouble," "The friendship is deteriorating," etc.). Others made reference to the nature of friendship in general (e.g., "Friendship is a two-way street," "Friendships take work," "Friendships can change for the better," etc.).

In the category pertaining to the message producer's own feelings, only one new item suggested itself. Several participants reported feeling nervous (scared, apprehensive, hesitant, worried, etc.) about the upcoming date with Terry. In the category pertaining to Terry's previous actions, many participants reported thoughts about the legitimacy of Terry's reasons for canceling (e.g. "I don't know if the cancelations were legitimate," "If he had a good reason, I wouldn't mind," etc.). In the category pertaining to Terry's traits, several novel traits were reported (e.g., trustworthy, organized, fake, tough, annoying, preoccupied,

[22] The Protocol-Specific/Situation-Specific coding system is reproduced as Appedix C.

spontaneous, disrespectful, etc.), but each occurred too rarely to merit inclusion on the checklist. In the category pertaining to Terry's motives, the only novel thought that occurred with any regularity had to do with Terry and another friend (e.g., "Terry has a new boy/girlfriend," "Terry wants to spend time with another friend," etc.).

One class of novel thoughts appeared in connection with the message producer's own actions: thoughts accepting responsibility for Terry's actions (e.g., "It's my fault," "I must have done something pretty bad to Terry," etc.). In the category pertaining to Terry's attitudes toward the message producer, one element that deserved inclusion had to do with the sincerity of Terry's commitment to the friendship (e.g., "Terry should want my friendship in a genuine way," "Terry should not feel like he has to be my friend," etc.). No new thoughts about Terry's behavior or Terry's thoughts emerged, but regarding future plans with Terry, several participants reported that they would call to confirm any future dates with Terry, and many reported thoughts about how to deal with Terry during the next encounter with him (most often, "I should confront Terry"). Table 1 summarizes my recommendations for additions to the thought checklist.

Table 1

Recommended Additions to the Thought Checklist

Topic	Specific Thoughts
The relationship	The friendship is in trouble. Friendships take work.
Own feelings	I am nervous about seeing Terry.
Terry's previous actions	I need to know if Terry's cancellations were legitimate.
Terry's motives	Terry wants to spend time with another friend.
Own actions	I must have done something to make Terry want to cancel.
Terry's attitudes	Terry should want my friendship in a genuine way.
Future plans with Terry	I should call to confirm dates with Terry. I should confront Terry.

General Discussion

This chapter began by asking the following three questions about the thought checklist: (a) Does the checklist task lead people to report more or fewer thoughts than they would in a concurrent ("on-line") elicitation technique? (b) Does the checklist itself provide people with substantively different ideas than they would have had on their own? (c) Is the checklist a reasonably comprehensive list of thoughts about the Broken Date Problem? We can now begin to address these questions. The investigation of the quantity of thoughts reported with different methods of elicitation showed that talk-aloud and thought checklist methods for thought elicitation yield essentially the same number of reported thoughts, and that fewer thoughts are reported in retrospective protocols than in either other condition. Next, since each category of thoughts on the checklist was well represented in the protocols, there is no reason to suspect that the checklist supplies people with ideas they would not have had on their own. Finally, the Checklist can be regarded as a reasonably comprehensive list of thoughts about the Broken Date Problem. Though a fairly large number of thoughts were classified as absent from the checklist, many of these either could not reasonably have been anticipated or were classified as absent only due to the stringent nature of the coding rules. No conceptually distinct categories of thoughts were reported. All in all, the thought checklist, with the recommended revisions, can be regarded as a valid measure of the thoughts people have as they read and respond to hypothetical communication scenarios.

Summary

This chapter reported a study comparing protocol analysis and thought checklist methods of thought elicitation. The average number of t-units reported across methods of thought elicitation was examined, and no significant effect was found. Subjects did report fewer t-units in retrospective reports than in talk-aloud protocols. The overlap between the methods was also examined by measuring the proportion of protocol-elicited thought-units that were synonymous with thoughts already on the thought checklist. Fifty-six percent of the protocol thought-units were found to overlap with thoughts on the thought checklist. This was taken as evidence of the checklist's representativeness. The residual thought-units were content analyzed and it was suggested that certain thoughts be included on subsequent versions of the thought checklist for the Broken Date problem.

CHAPTER SIX
DEVELOPMENT AND VALIDATION OF THE MESSAGE-FRAME CODING SYSTEM

The model of message production being developed here seeks to capture the relationship between thoughts and expressed message contents. The thought checklist discussed in chapter 5 was developed so participants' thoughts could be represented in the PDP model as vectors of numeric values. Before actually building and training the model, it is necessary to develop a system for coding messages into a numeric vector representation. This chapter describes the development and validation of a system for coding the elements of Broken Date messages into a numeric vector representation.

Overview

This chapter is organized into four sections. The first section discusses the theoretical and methodological rationale for the construction of the coding system. Two primary concerns in this regard were (a) to distribute the representation of message contents and (b) to reliably assign a variety of message elements to a large number of categories. The second section describes specific steps taken in the development of coding rules and categories. The third section describes the implementation of the coding system as a computer program. An attempt to automate the coding process was undertaken to address reliability concerns that arise in a system with many categories. The fourth section assesses the performance of the automated version of the coding system on novel messages.

Representation of the Broken Date Messages as an Activation Vector

The Broken Date messages were segmented into sets of message elements. A message-element is any independent clause; phrases that constituted completed thoughts (e.g., Okay, Terry) but that did not contain both a subject and a verb were nonetheless counted as separate elements.

A classification system was developed in order to group synonymous message elements. This system contained 72 categories. Message elements were taken to express the same idea when they differed only in using alternative terms to refer to a particular feature of the situation or when they expressed the same idea within alternative syntactic structures using common vocabulary. The number of required categories was further reduced by specifying idea-complexes, and allowing each of several different idea-types to represent all or part of the idea-complex. These idea-complexes are henceforth referred to as "message frames" or simply "frames". The 72 message frames are reproduced in Table 2.

Redundancy Within Message Frames

As synonymous message elements were grouped together, a surprisingly high degree of parallelism and redundancy became apparent. People seemed to draw on a fairly large but finite set of stock phrases to address the different topic areas that were relevant. Granted, different people touched on different topics (e.g., some dwelt on the friendship, some on Terry's bad behavior, some on the request to reschedule, etc.), but when two or more people addressed the same topic, they did so using remarkably similar phraseology (see Table 3).

Becker (1975) has argued that much of language production is accomplished by the systematic exploita-tion of a phrasal lexicon. The phrasal lexicon (an analogy to the mental lexicon of words, see Levelt, 1989, esp. chapter 6) is a database of preformed phrasal structures that can be drawn upon to express common ideas. Language production with the phrasal lexicon involves the retrieval and sequential organization of fairly high-level structures (phrases). Becker's model differs from standard accounts of production in claiming that the primitive units are already-formulated phrases. Standard models claim that

Table 2

Frequency and Classification Data for 72 Message Frames

Frame	No. Present		Mean Freq.[c]	Test Proportion Correctly Classified
	Orig.[a]	Test[b]		
You have broken our dates a lot lately.	69	47	.362	.128
Sorry, I can't go tomorrow.	59	45	.325	.289
Call me when you have more time.	52	35	.272	.171
Okay, I can go tomorrow.	61	24	.266	.500
I'm tired of your behavior.	33	32	.203	.250
Your behavior is bothering me.	39	23	.193	.043
What is going on?	29	32	.191	.188
I understand your constraints.	40	19	.184	.105
We have been good friends for some time.	31	24	.172	.583
You've been busy.	29	22	.159	.455
If you don't want tosee me, just say so.	21	27	.150	.111
I want to get together.	32	15	.147	.467
Is there a problem?	24	19	.134	.316
Check with me another time.	10	27	.116	.000
When we make dates make sure you are free.	13	23	.113	.000
Tomorrow is fine if we can make it.	24	12	.113	.083
Your behavior is unacceptable.	9	26	.109	.038
Since you cancel, I won't make plans with you.	16	19	.109	.000
I expect you to cancel tomorrow, too.	31	4	.109	.250
If our friendship is impor tant, then make time for me.	25	10	.109	.000
Goodbye.	14	16	.094	.438
Forget it, Terry.	15	15	.094	.200
I'll find alternatives.	18	12	.094	.000
I have a busy life, too.	10	17	.084	.176
You leave me with nothing to do.	21	6	.084	.000
Keep our date tomorrow.	20	6	.081	.167
We need to talk.	17	9	.081	.000

Table 2 (cont'd.)

Frequency and Classification Data for 72 Message Frames

Frame	No. Present		Mean Freq.[c]	Test Proportion Correctly Classified
	Orig.[a]	Test[b]		
You have made me a low priority.	16	10	.081	.000
I want to see the movie tonight.	11	14	.078	.071
I suspect that you are avoiding me.	20	4	.075	.500
What is the story on the meeting.	17	7	.075	.143
Our friendship is suffering.	15	6	.066	.000
You should consider your priorities.	12	8	.063	.125
Okay, go ahead and cancel.	10	10	.063	.100
Let's find a compromise plan.	15	4	.059	.000
Our friendship may end.	15	3	.056	.000
Friendship takes mutual effort.	12	6	.056	.000
Let's talk tomorrow about the date.	10	8	.056	.000
Understand my perspective.	9	8	.053	.000
You don't have time for me.	8	8	.050	.250
If you don't want dates don't make dates.	9	7	.050	.143
I want to remain your friend.	8	7	.047	.143
I have tried to set up dates with you.	11	4	.047	.000
You have not been a good friend lately.	12	3	.047	.000
I've made sacrifices to be with you.	9	6	.047	.000
If you called sooner I could have replaced you.	8	7	.047	.000
If you cancel again, I'll be angry.	10	5	.047	.000
I miss you.	10	3	.041	.333
Psychological prefaces	3	10	.041	.200
I make time for you.	9	3	.036	.000
I'm not angry.	5	6	.034	.167
Our friendship is important to me.	4	7	.034	.143
We haven't been together for a while.	9	2	.034	.000
We have done so much together.	6	3	.028	.000

Table 2 (cont'd.)

Frequency and Classification Data for 72 Message Frames

Frame	No. Present		Mean Freq.[c]	Test Proportion Correctly Classified
	Orig.[a]	Test[b]		
If our friendship isn't im portant, then forget it.	7	2	.028	.000
Skip the meeting and keep our date.	6	3	.028	.000
I care for you.	7	1	.025	.000
We had a date.	5	3	.025	.333
Be more organized.	7	1	.025	.000
I want to have an honest talk.	5	2	.022	.000
I'll stop bothering you.	7	0	.022	-
You take me for granted.	6	1	.022	.000
Friends are important.	3	3	.019	.000
I can't count on you.	3	3	.019	.000
I hope we get together tomorrow.	4	2	.019	.000
You need to make choices.	4	1	.016	.000
Maybe I'm overreacting.	4	1	.016	.000
We can solve this problem.	4	1	.016	.000
Do not feel obliged to me.	1	4	.016	.000
You should make some compromises.	3	1	.013	.000
I don't know what to think.	1	3	.013	.000
I'm sorry for my behavior.	2	0	.006	-

[a]Number of instances of this frame from O'Keefe and Lambert (1989) data.
[b]Number of instances of this frame from chapter five Thought Checklist condition.
[c]Cumulative mean frequency per subject (N = 320).

Table 3

Selected Message Frames Illustrating Within-Frame Redundancy

Frame #2: I'm tired of your behavior.

I am just tired of being stood up and canceled out by you.
I am really getting tired of you postponing all of our plans.
I am very fed up with the cancelations you have been making lately.
I would like you to know I am really getting tired of your backing out on all of our plans.
I am getting kind of sick of you always canceling out.
I am sick of you blowing me off.
I am really sick of being blown off.
Terry I am sick of the way you have been treating me lately.
I am getting sick of your bs.
Terry I am getting sick and tired of you ruining all of our plans all the time.
Listen Terry I am really getting fed up with your breaking plans.
Terry I am sick and tired of you canceling on me.
I am sick of it.
I do not need this.
I am sick of your excuses.
I am sick and tired of your meetings.
I am sick of working around your schedule.
I do not need that.
Terry I am tired of having you cancel out all the time.
I am sick of being blown off.
I am really getting tired of this Terry.
I am really hurt.
I am tired of being disappointed.
Terry I am really tired of your games.
I am really sick of making plans with you and having them broken all the time.
I am really tired of you canceling dates with me.
I am sick of this.
I am getting tired of your constantly breaking our date.
I am tired of getting canceled on.
Enough is enough.
I can not handle much more of this.
I just want to say I am sick and tired of being your little puppy.
I am tired of making tentative plans with you and having you cancel them.
I am sick of it.

Frame #5: What is going on?

What is up with you?
I do not know what your problem is.
Your other commitments cannot take up all of your time.
Terry I do not understand.
Which is it?
Do you want me to stop begging you?

Table 3 (cont'd)

Selected Message Frames Illustrating Within-Frame Redundancy

What is going on lately anyway?
What is the deal?
Do you just not want to go?
I do not know what the problem is.
What has happened?
It is hard to understand your recent business.
Please let me know what is happening.
What are you doing that is so important that you can not do anything with me anymore?
What is the deal?
What is happening?
What the hell is going on?
What is the deal?
What is the matter?
I just do not understand what is going on here.
You are not making any sense to me.
I could understand if something popped up once in a while.
What is up?
It is all right if you have to cancel a few of our engagements.
What is going on?
What is your problem?
What is up with our relationship?
Tell me what is going on.
I want to know what is going on?

Frame #6: Is there a problem?

Tell me what is wrong.
Have I done something to make you want to avoid me?
Is there something wrong?
Are you mad at me?
I am not boring or something am I?
Is anything wrong?
Is there something wrong?
Are you mad at me about something?
Is there something wrong?
Is something wrong?
Is this the case?
Did I do something?
Is there something wrong that is bothering you?
Is there anything I can do?
Is there another guy?
Are you unhappy with our times together?
Are you mad at me?
Is something wrong?
Have I done something to make you mad?
Why can not we just get together?

Table 3 (cont'd)

Selected Message Frames Illustrating Within-Frame Redundancy

Are there any problems I should know about?
Is it something I did or said?
Have I done something wrong?
Is it something I did?
Is there something other than the meeting keeping you from going out?

Frame #12: We have done so much together.

We have done a lot together.
We used to do so much together.
I always thought we were having fun.
We used to spend a great deal of time together.
We used to go to a lot of places together.
We have done everything together.

Frame #23: We need to talk.

I really want to talk about this.
If we are growing apart we have to decide what the best move for each of us is.
We need to talk.
We have lots of catching up to do.
We need to talk.
Terry I have been meaning to have a talk with you.
I think we should talk.
Terry I think we need to talk and get some things straightened out.
Terry we really have to talk.
I wish you would talk to me about the problem.
I think we have to talk.
I have to talk with you.
It is about time we discussed the problem that has developed in our friendship.
I have really been wanting to talk to you lately.
Our next date is going to be a long conversation.
We really need to talk about your behavior.
We need to talk about why you have not had time for me lately.

Frame #50: Let's talk tomorrow about the date.

Call me early to confirm our date.
Why do not you give me a call tomorrow night.
I will call you to confirm our date.
If you actually can make the movie give me a call tomorrow.
Why do not you call me tomorrow and let me know how your agenda looks.
Call me tomorrow.
Give me a call tomorrow night.

Table 3 (cont'd)

Selected Message Frames Illustrating Within-Frame Redundancy

Call me later to confirm tomorrows date.
Give me a call tomorrow night to make sure I am free.
Call me tomorrow.
Just call me tomorrow.

Frame #53: I understand your constraints.

I understand you miss dates.
I understand that you have other things to do.
I have tried to understand.
I understand you have obligations and responsibilities.
Okay Terry I understand.
If you have things to do you have things to do.
I understand that you are busy or whatever.
I understand if you are too busy to go out.
I guess meetings are pretty important.
I have tried to be understanding about your hectic schedule.
I understand you have a lot to do.
I understand if you have a meeting.
I realize you have to go places and do things.
I understand that you have other commitments.
I understand that things do not always work out.
I could understand if you were out with a gal.
Terry I realize you have a lot of other obligations.
I realize we are both very busy.
I understand.
I will understand.
I am not saying you cancel for the wrong reasons.
I realize you have a busy schedule.
I have been trying to understand that you have been busy lately.
If you honestly think you will not be able to go tomorrow thenyou do not have to take me.
I realize that your activities are very important to you.
I understand that you have a meeting to go to today.
I accept that you have a meeting tonight.
I know you have a meeting tonight.
I understand your other obligations.
I realize this meeting is important.
Your meetings and group activities are important to you.
If you have to go to meetings fine.
I understand that you have meetings and everything.
I can understand that you have other things in your life.
If you really do have a meeting then I guess I understand.
I know you have lots to do.
I understand that there are a lot of worthless meetings we have to attend at school.
I understand this meeting arose suddenly.
I understand about the meeting.

Table 3 (cont'd.)

I understand the meeting.
I understand.

---

surface structures are realized by processes that operate on lower-level entities (e.g., phonemes, morphemes, words, see esp. chapter 5 of Levelt, 1989) and that these low-level objects are systematically fit into abstract structural frames (cf. Dell & Juliano, 1990). The data from this study seem to provide strong support for a model of production based on something like a phrasal lexicon. It is argued that the set of frames discussed here comprises the Broken Date phrasal lexicon.

Reliability

Ordinary procedures for establishing the reliability and consistency of a coding system are not well designed to addressthe problems created by a 72-category coding system. Among other things, it is exceedingly difficult to keep 72 different message frames in mind; human coders have difficulty in applying such a system consistently in a reasonable amount of time. In addition, standard tests of reliability (e.g., Cohen, 1960) assume that instances of each category will occur frequently in each reliability sample. But with so many categories, even relatively large reliability samples may not contain any instances of low-frequency categories. Consequently, to address the need to establish that the system is consistent and to ensure that it was reliably applied, a prototype automatic coding system dubbed AutoCoder (Lambert, 1991) was developed and implemented as a LISP computer program.

Related work on automated coding. When confronted with the same reliability dilemma, previous researchers have also taken steps to automate the coding process. The most extensive automatic coding system discussed in the literature was developed to automate the coding of think-aloud and talk-aloud protocols (Waterman & Newell, 1971, 1972, 1973). Waterman and Newell's (1971) protocol analysis system (PAS-I)[23] operated on protocol fragments already segmented into topic segments by the researchers. Topic segments were independent clauses, similar to the message elements discussed above. PAS-I "used a keyword grammar to map the protocol segments onto propositional elements. This grammar can be seen as a formal specification of the set of linguistic patterns that should be mapped on a given propositional relation" (Ericsson & Simon, 1984, p. 301). Further discussion and evaluation of PAS-I are available in the literature cited. For now, it is important to note that there is a precedent for automating the coding process, and that this precedent involves using a keyword grammar to translate segments of text through several levels of abstraction.

The next section describes in more detail the steps that were taken to implement AutoCoder, a keyword-based automatic coding system.

AutoCoder: Development of a Prototype Automated Coding System

AutoCoder is a simple pattern matching algorithm that classifies messages at several levels of abstraction. There are five levels of abstraction to consider in AutoCoder: (a) messages (b) message elements, (c) keywords, (d) phrases, and (e) frames.

Messages

External to AutoCoder are the actual messages as they appear on study participants' questionnaires. Below is the text of one subjects message:

No, I will be busy tomorrow night. How do all of these things suddenly pop up out of

[23] There was also a PAS-II (Waterman & Newell, 1972, 1973).

nowhere? Don't you have a set calendar of events? I could understand if something popped up once in a while but three times in one week is ridiculous. (O'Keefe & Lambert, 1989, #156)

Message Elements

The most basic representation of messages within AutoCoder is the message-element, represented as a list of words. Each Broken Date message was segmented into message elements as described above. In addition, all punctuation was removed; conjunctions occurring between independent clauses were removed; contractions were expanded; and spellings were regularized. Thus the above message would appear as follows:[24]

```
((No I will be busy tomorrow night)
(How do all these things suddenly pop up out of nowhere)
(Do not you have a set calendar of events)
(I could understand if something popped up once in a
 while)
(Three times in one week is ridiculous))
```

Keywords

At the next highest level of abstraction are keywords. Keywords are defined as sequences of words. A single keyword definition may consist of several distinct sequences of words, each of which is an instance of the keyword. For example, the keyword *BUSY* is represented within AutoCoder as follows:

```
(*BUSY*
    ((simply) busy (and stuff))
    (getting in over your head)
    (involved (in activities))
    (hectic)
    (involved in other things)
    (involved in *AN-AMOUNT* *THINGS*)
    ((very) active (now))
    (a very busy person)
    (booked up))
```

Two additional aspects of keyword definitions are noteworthy. First, some words or sequences of words in keyword definitions may be optional. Optional sequences are enclosed within parentheses. In the example given, both "active" and "very active" would be classified as *BUSY*. Second, keywords can be defined in terms of other keywords. The treatment of such "compound" keyword definitions is dealt with below as part of the discussion of the AutoCoder algorithm.

Phrases

At the next highest level of abstraction are phrases. Phrases are defined by lists of keywords. Some keywords may be optional. Phrases may not be defined in terms of other phrases. Below is the phrase definition that would match the first element in the example message:

```
(*54D* (*NO*) *I* (*IS*) (*ALREADY*) (*REMAIN*) *BUSY*
    (*TOMORROW*) (*TERRY*))
```

The string "`*54D*`" means that this phrase is member of the set of phrases that define frame 54 ("Sorry, I can't go tomorrow").

Frames

At the highest level of abstraction are message frames. message frames are defined in

[24] Extracts from AutoCoder source code or from AutoCoder data structures appear in a different font.

terms of sets of phrase definitions. So for example, frame number 67 ("Be more organized") is defined by the following phrases:

(*67A* (*YOU* *SHOULD*) *BE-RESPONSIBLE*)
(*67B* *FAIL-TO* *YOU* *HAVE* *A-CALENDAR*)
(*67C* (*PLEASE*) (*IN-THE-FUTURE*) *RECORD-OUR-DATE*)
(*67D* *IF* *YOU* *HAVE* *A-CALENDAR* *YOU* (*WILL* *HAVE*)
UNDERSTAND *THAT* *YOU* *HAVE* *A-MEETING*)
(*67G* *DID* *YOU* *LOSE* *A-CALENDAR*)
(*67H* *FIND-OUT* *ABOUT* *A-MEETING* *SOONER*)
(*67J* *I* *WANT* *YOU* *TO* *BE-RESPONSIBLE*)

In Table 4, one can see each of the elements of message 156 represented at three levels of abstraction.

Table 4

Message Elements at Three Levels of Abstraction

Message element:	(No I will be busy tomorrow night)
Phrase:	(*NO* *I* *REMAIN* *BUSY* *TOMORROW*)
Frame:	"Sorry, I can't go tomorrow"
Message element:	(How do all these things suddenly pop up out of nowhere)
Phrase:	(*HOW* *DO* *AN-AMOUNT* *THINGS* *ARISE*)
Frame:	"What is the story on the meeting"
Message element:	(Do not you have a set calendar of events)
Phrase:	(*FAIL-TO* *YOU* *HAVE* *A-CALENDAR*)
Frame:	"Be more organized"
Message element:	(I could understand if something popped up once in a while)
Phrase:	(*I* *UNDERSTAND* *IF* *YOU-CANCEL-A-FEW-DATES*)
Frame:	"What is going on?"
Message element:	(Three times in one week is ridiculous)
Phrase:	(*YOUR-BEHAVIOR* *IS* *UNACCEPTABLE*)
Frame:	"Your behavior is unacceptable"

The AutoCoder Algorithm

The AutoCoder algorithm is a simple pattern matching algorithm that transforms message elements into lists of keywords, lists of keywords into phrases, and phrases into frames. For a pattern at one level of abstraction to be successfully classified into the next highest level, each element of the candidate pattern must exactly match the target pattern at the next highest level (excepting optional elements). There are no variables, and there is no element of probability in classification. The match between a candidate pattern and a target pattern is all-or-nothing. The input to AutoCoder is a list of messages. Each message is a list of message elements, and each message-element is a list of words.

Words-to-Keywords

The first step in coding a message-element is to transform a list of words into a list of keywords. This is done by searching the list of keyword definitions for a definition that matches the current list of words. At first the list of words is the entire message-element. If no keyword definition matches the entire message-element, then the last word in the message-element is pruned off the tail, and the list of definitions is searched again with the shortened word list as the pattern to match. This procedure is repeated until a match occurs or until there is only one word remaining from the original word list. If it matches, the matched keyword is appended to a list of keywords, the words that match are pruned from the front of the keylist, and the process is repeated on the remaining list of words. If a single word fails to match any keyword definitions, the word is not defined, and partial failure is reported. The word itself is pruned from the head of the list of words, appended to the list of keywords, and the process begins again. Pseudo-code for the WORD-TO-KEYWORD classification function is given in Figure 4.

```
WORD-TO-KEYWORD (list-of-words)
    UNTIL list-of-words is empty
     IF list-of-words contains one word
      AND word fails to match any keyword-definition
            THEN return word and report failure and call WORD-TO-KEYWORD on  rest of list-of-words
            ELSE
       IF list-of-words matches some keyword-definition
            THEN add the matching keyword to the keyword-list
             AND call WORD-TO-KEYWORD on the rest of the list-of-words
            ELSE prune the last word from the list-of-words
             AND call WORD-TO-KEYWORD on the shortened list-of-words
    RETURN keyword-list
```

Figure 4. Pseudo-code for WORD-TO-KEYWORD classification.

In general, the search strategy proceeds in a left-to-right manner, preferring to classify the longest sequences first. For example, if the message-element being examined were "you always cancel" the system would try to match the following subsequences, in order:

```
(you always cancel)
(you always)
(you)
(always cancel)
(always)
(cancel)
```

The keyword definitions are stored in a discrimination net to facilitate parallel search. A discrimination net is a tree-like data structure whose root corresponds to the beginning of a list of words to be matched. The tree has one branch for every keyword definition. Search of the definitions proceeds by searching along each branch until a word in the list fails to match or until every word has matched and the end of a branch has been reached.

Keywords-to-Phrases and Phrases-to-Frames

Once a message-element is successfully coded into a list of keywords, the list of phrase definitions is searched for a phrase definition that matches the current list of keywords. The search proceeds in essentially the same fashion as described above, but the data structures are now one step higher in the abstraction hierarchy.

One important difference in the classification of keyword-lists into phrases concerns the presence of psychological prefaces. A psychological preface is a special kind of message-frame. These elements (e.g., "I believe," "I think," "It appears," "It is like," "It seems like," "I feel that," etc.) act as interchangeable place-holders in the stream of expressed content. Before trying to classify a keylist as a phrase, AutoCoder searches for psychological prefaces. If one is present (at the front of a list of keywords) it is pruned off, and classification proceeds on the remaining keywords in the list.

The large number of optional elements in some phrase definitions makes it impractical to expand each option into a distinct path as was done for keyword definitions, so phrase definitions are searched serially rather than in parallel. Once a list of keywords is classified as a phrase (e.g., as phrase *54D*), it is implicit that the phrase is part of the definition of frame 54.

Quantitative Characteristics of Keyword and Phrase Definitions

The coding system was constructed using 207 Broken Date messages originally collected for O'Keefe & Lambert's (1989) study of the thought checklist. There were 5174 total words in this corpus, but only 773 unique words. There were 1127 message elements in the corpus, an average of 5.44 elements per message.

Looking at the degree of data compression achieved by the coding system is helpful. For example, there were 477 keywords defined in the system, 61.7% of the number of different words in the corpus. There were 666 phrases defined in the system, 59.1% of the total number of message elements. Similarly, there were an average of 4.93 keywords per phrase, 54.6% of the average number of words per message-element. These comparisons illustrate that the coding system achieved a substantial measure of abstraction from the instances on which it was based.

Performance on Unseen Cases

One important measure of the quality of the coding system is its success at classifying novel instances. In the final analysis, AutoCoder will be a solution to the reliability question only if it aids in the coding of new instances. To test AutoCoder's generalization performance, 113 Broken Date messages (a total of 796 message elements, taken from the thought checklist condition of the protocol experiment from chapter 5) were input to AutoCoder. AutoCoder correctly classified 127 or 15.9% of these novel elements (correctness in coding was assessed by comparison to the author's manual codings).[25] In absolute terms the success rate is low, but it is not clear how this measure of performance should influence the overall evaluation of AutoCoder.

Two factors make it especially difficult to assess the system based on this number alone. First of all, there is no other automatic coding system whose performance can be

[25] The entire corpus of message elements, sorted by message frame, is included as Appendix D. The correctness of the final classifications was verified by the author.

compared directly to AutoCoder's. However, an indirect comparison can be made to PAS-I. Ericsson and Simon (1984, p. 305) report a 62% correct classification rate for PAS-I. The reliability of this performance estimate should be qualified, as it is based only on 100 segments (38 of which were either completely or partially mishandled).[26] Also, the protocols used to test PAS-I came from a cryptarithmetic task, a task with a phrasal lexicon arguably much narrower than the Broken Date task. Even so, PAS-I outperformed AutoCoder on this small sample. Second, granted that AutoCoder is a prototype system relying on exact matching and possessing no explicit "knowledge" of inflection, verb conjugation, tense, or syntax, it is not realistic to expect much better generalization performance than has been observed. Before passing final judgment on the prototype system's performance, the nature of the classification successes and failures will be examined in more detail.

Word-to-keyword misclassifications. Overall, AutoCoder was very successful at coding words and clusters of words into keywords. Of 5713 total words in the corpus of novel message elements, 4820 or 84.4% were correctly classified as keywords. The words that were not classified were often idiomatic (e.g., be straight [with me], running around [on me], left in the dust, two-way street, play games [with me], backing [out on me], night after night, kiss off, hanging [by] a string, pulling away, for God's sake, grain [of] salt, taste [of your] own medicine, etc.). Some failures were common words from the corpus that had not appeared alone in the earlier sample. That is, words like the or cancellations, had previously appeared in keyword definitions only as part of larger expressions (e.g., "your frequent cancellations," "the movie," etc.).

Keyword-to-phrase misclassifications. The bulk of the misclassifications resulted from the system's inability to deal with any unanticipated variation in keyword order. Table 5 displays three examples of keyword-to-phrase misclassification.

Table 5

Examples of Keyword-to-Phrase Misclassification

Message-element:	(I really miss you Terry)
Keywords:	`(*I* *REALLY* *MISS* *YOU* *TERRY*)`
Phrase Definition:	`(*14A* *I* (*REALLY*) *MISS* (*YOU*) (*SEEING-YOU*))`
Message-element:	(Have I done something to upset you)
Keywords:	`(*DID-I-DO* *SOMETHING* *TO* *BOTHERED* *YOU*)`
Phrase Definition:	`(*6J* *DID-I-DO* *SOMETHING-WRONG*)`
Message-element:	(What is going on Terry)
Keywords:	`(*WHAT* *IS* *HAPPENING* *TERRY*)`
Phrase Definition:	`(*5D* (*TELL-ME*) *WHAT* (*THE-HELL*) *IS* *HAPPENING* (*HERE*) (*LATELY*))`

[26] Note that AutoCoder was tested on a corpus nearly eight times this size (796 elements).

It is clear from the examples in Table 5 that the misclassifications were often due to minor mismatches. In two cases, the phrase definition simply lacked the appellation "Terry" at the end of the phrase. In the other case, the word "upset" was classified as belonging to the keyword *BOTHERED*. If it had been classified as *BOTHER*, then the set of keywords *TO* *BOTHER* *YOU* would have bben classified as the compound keyword *SOMETHING-WRONG*, and the phrase definition would have successfully matched. This is an instance where the past tense of the irregular verb "upset" caused a problem for AutoCoder.[27]

From the raw classification statistics, it is not clear to what the poor generalization can be attributed. There are two obvious possibilities. First, the system could be too simple. It simply might not be possible for an "unintelligent" pattern matching algorithm to correctly classify novel linguistic expressions. Claims about the creativity and generativity of language lend credence to this explanation. If this first explanation turned out to be accurate, then AutoCoder must be deemed inadequate in principle.

The second explanation is not as serious, and points to a practical problem rather than a problem in principle. The poor generalization might be due to the limited number of examples on which the coding rules were based. Even if the second explanation turns out to be accurate, it does not spell unqualified success for AutoCoder. Rather, it means that the idea should be tested further with a larger example set.

Analysis of misclassifications. If the second explanation were true, and poor generalization were due to the fact that relatively rare message frames were not well captured in the current coding rules, then one would expect better generalization performance on more common frames. This hypothesis was tested by comparing the proportion of correctly classified test instances of each frame to the number of instances of each frame in the corpus from which AutoCoder was built. In order to compute the proportion of novel elements classified into each frame, it was necessary to hand code the remaining 669 unclassified message elements into the appropriate message frames. This was accomplished by comparing the partially classified elements (either lists of keywords or hybrid lists of words and keywords) to the instances already covered by AutoCoder.

Frames with many instances in the original corpus were expected to show better generalization performance. The number of instances of each frame in both the original and the testing sample as well as the proportion of testing instances correctly classified into each frame are given in Table 2. The proportions underwent an arcsin transformation to insure normality before being entered into the analysis (Cohen & Cohen, 1983). There was a significant positive relationship between the number of instances of each frame in the original corpus and the proportion of novel instances correctly classified into each frame ($r(68) = .467$, $p < .0001$).[28]

Discussion of misclassifications. The discovery of a significant relationship between the number of instances of a frame in the original sample and the generalization performance lends qualified support to the second, less severe explanation of AutoCoder's poor generalization performance. Given this result, it seems likely that generalization performance will improve as new coding rules are added to cover more instances of less frequently occurring frames. This result is not surprising given that AutoCoder is an exemplar-based classification system, and that the generalization performance of exemplar-based classifiers typically increases as a function of the number of examples stored.

[27] Tense is not handled consistently in AutoCoder, and this is an obvious weakness.

[28] There were 68 rather than 70 degrees of freedom in this test because two frames had no instances in the testing set. Proportions of correct classifictions could not be calculated for these frames.

AutoCoder, or systems like it, will prove to be useful if the generalization performance reaches a satisfactory level before the number of stored-exemplars becomes unmanageably large.

Possible extensions. Though adding more coding rules to cover less frequently occurring frames is likely to improve performance, more fundamental extensions of the system suggest themselves as well. One specific extension is described. As it stands, AutoCoder is like a glorified look-up table. An element to be matched is compared to all the elements AutoCoder has stored, and if it matches one exactly (excepting options), it is classified.[29] Exact matching is probably too stringent a criterion for classification in this context. Classification should be made slightly probabilistic. For example, if 4 out of 5 keywords in a given pattern match a stored phrase, the list of keywords could be classified as matching that frame with probability equal to .80. As the system stands, such a pattern would be rejected as a failure to match. This extension is in the process of being implemented. Others, including variable binding, should be considered.

Summary

This chapter addressed the methodological problem of transforming messages into a numeric vector representation so they could be input to machine learning algorithms. The development of a classification scheme for message elements (independent clauses) was described. Similar message elements were grouped into idea-complexes called message frames. The high degree of redundancy and parallel structure within frames was noted and discussed in relation to Becker's (1975) idea of a phrasal lexicon.

In the second half of the chapter, AutoCoder, a computerized pattern matching algorithm for coding messages into message frames, was described. The system was able to correctly classify approximately 16% of the novel message elements it attempted. The underlying cause of this relatively poor generalization performance was analyzed. There was a moderately strong positive relationship between the number of instances of each message frame in the original corpus and the proportion of novel instances correctly classified into each frame. Based on this result, it was suggested that generalization performance could be improved by adding more exemplars to AutoCoder. An extension to the system that would allow it to match probabilistically was also described.

Though the problem of reliably coding thousands of instances into 72 categories has still not been completely solved, the prototype AutoCoder system was offered as evidence that the problems are not insurmountable in principle. In fact, Ericsson and Simon (1984) suggest that automated coding systems need not be perfect to fulfill their function. Even partially successful automated systems and semi-automated systems (e.g., Bhaskar & Simon, 1977) increase the reliability of coding while substantially easing its difficulty in schemes with many instances and categories. Judged against this criterion, the construction of AutoCoder is seen as a qualified success.

[29] The analogy to a look-up table does not acknowledge that instances are classified through several layers of abstraction.

CHAPTER SEVEN
A CONNECTIONIST MODEL OF CONTENT SELECTION IN MESSAGE DESIGN

The purpose of this chapter is to determine whether a PDP model mapping patterns of thoughts onto patterns of message- frames has face validity, and if so, to begin to illustrate an inductive approach to the study of distributed goals, situated logics, and other concepts essential to a situated theory of message design.

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 this model: (a) no explicit logic of action intervenes between goals and messages; (b) message types play no explicit role in processing; and (c) goals are represented as distributed patterns of thoughts. Implementing these design features in a PDP framework may avoid some of the problems that confront planning models.

Several research questions motivate the construction of the thought-to-message simulation: (a) Can PDP pattern association networks capture relations between patterns of thoughts and message-frames? (b) Can PDP modelling be used as a kind of exploratory data analysis technique (e.g., Hanson & Burr, 1990) for describing the relations between thoughts and message-frames? (c) Can goals be modelled as patterns of thoughts? (d) Is it sensible to conceive of a message design logic (O'Keefe, 1988) as being embodied in learned patterns of connections between thoughts and expressed message elements?

Overview

This chapter is developed in several sections. The first section reviews other work on language production that has influenced the current model. The second section briefly reviews the structure and functioning of the current model at a theoretical level. The third section describes the methodology used in constructing, training, and evaluating the PDP model. Included here are descriptions of the example representation, learning algorithm, and evaluation criteria. A select subset of the weights from a trained network is analyzed and interpreted in the fourth section. The fifth section argues, based on cluster analysis, that distributed patterns of thoughts function as goals. The sixth section describes the potential uses of the proposed model in the study of individual differences in communication competence, and the final section is a summary.

Related Models of Language Processing

PDP Models

In recent years, several PDP models of language processing have appeared in the literature. There has been a flurry of work in this area because, for many cognitive scientists, the merit of the PDP framework in general is measured by its success in accounting for linguistic phenomena in particular. But since PDP models are based on networks of associations between structureless primitive elements, many have argued that they are fundamentally maladapted for modelling language and cognition.

A well-known PDP model of the acquisition of the past tense of English verbs (Rumelhart & McClelland, 1986) served simultaneously as the inspiration for subsequent PDP models of language processing and as the object of vehement criticism (see e.g., Fodor & Pylyshyn, 1988; Lachter & Bever, 1988; Pinker & Prince, 1988). While some of the criticism levelled by Pinker and others may be valid, the pace of PDP modelling of linguistic phenomena has not slowed. Several of Pinker, et al.'s objections are addressed in a model of the acquisition of the definite article in German (MacWhinney, Leinbach, Taraban, & McDonald, 1989). At the sentence level, case-role assignment (McClelland & Kawamoto,

1986) and grammatical parsing (Hanson & Kegl, 1987) have been modelled. The representation of lexical knowledge and lexical categories has also been modelled (Elman, 1990; Ritter & Kohonen, 1989). Further, at the level of phonological processing, a substantial body of work exists (see, e.g., Dell, 1986, 1988; Dell & Juliano, 1990; Dell & Reich, 1981). This is not meant as an exhaustive review, rather as an illustration of the diversity of linguistic phenomena that have been modelled in the PDP framework.

The choice here of the PDP modelling framework is influenced by Dell and Juliano's (1991) discussion of how PDP models blur the structure/content distinction so prevalent in linguistic theorizing and by Elman's (1990) discussion of how PDP models provide a new way of capturing type/token relations. Since much of the earlier critique of plan-based models was found to hinge on a commitment to untenable form/function and type/token distinctions, and since PDP models claim to provide a new angle on these distinctions, it seems sensible to adopt the PDP modelling strategy.

Local Management Models

In addition to the PDP framework, this work is also influenced by local management conceptions of planning (Agre & Chapman, 1987; 1990) and language production (Sibun, 1991b). The local management image of planning derives 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).[30] 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 present model differs in its details from the kind of cognitive machinery Agre and Chapman suggest, both were built in an effort to explain the cognitive substrate that supports improvisation during situated activity.

Finally, Sibun's (1991a, 1991b) Salix, a system for generating natural language descriptions of houses and families, embodies design principles and theoretical commitments that are shared by the model presented here. Most notably, Salix generates coherent text without using abstract structural representations and without reference to act types or 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 from among a small set of locally available alternatives what to say next. In a similar vein, the next section introduces a PDP model for locally organizing elements of content into coherent messages, without reference to act types or hierarchical plans.

Review of Theoretical Model

This section summarizes the basic principles that constitute the proposed PDP model of message production. The model associates patterns of thoughts with patterns of message-frames (see Figure 5). The mapping is analogous to Herrmann's (1983) proposed mapping of the propositional base of an utterance onto its semantic input (see chapter 4). Furthermore, the set of message-frames identified in chapter 6 can be construed as the Broken Date phrasal lexicon (Becker, 1975), that set of content elements that are potentially relevant for expression in the social situation described by the Broken Date scenario. The items in the phrasal lexicon are indexed (or content-addressed) by patterns of thoughts, where patterns of thoughts play the role normally reserved in plan-based models for discrete goals.

[30] Compare this to the definitions of language production reviewed in chapter 2.

In plan-based models, production begins with a discrete goal and proceeds by the construction and execution of a hierarchical plan to achieve that goal. Natural language elements are produced when the abstract linguistic entities identified in the plan are instantiated as concrete linguistic expressions. In the PDP model, production begins with a goal (represented in distributed fashion as a pattern of thoughts) and proceeds by passing activation to message-frames. The message-frames that receive activation in excess of a given threshold are output by the system. There are no plans. There is no functional abstraction and no instantiation of abstract types.[31]

Methods

PDP pattern association networks learn their connection weights from examples. The precise nature of this learning process is described in more depth below. For now it is enough to note that the model learns from "correct" examples of the function it is trying to approximate.[32] In order to train the thought-to-message simulation, it is necessary to gather examples of "correct" thought-pattern/message-pattern pairs. The methods for thought and message elicitation are described in chapters 5 and 6 respectively. The input features for this simulation are precisely the same as those used in the thought checklist from chapter 5. The output features correspond to the message-frames described in chapter 6.

Subjects

The data for this model came from the n = 207 subjects from O'Keefe and Lambert (1989) and the n = 113 subject from the thought checklist condition of chapter 5. Thus the models are based on a total N = 320.

Representation of Input Features

Each subject completed the Thought Checklist task for the Broken Date scenario as described in chapter 5. Subjects marked an "X" next to each thought that they recalled having as they responded to the Broken Date scenario and left the others blank. For each subject, a pattern of thoughts was represented as a vector of binary values on the 121 variables constituted by the thoughts on the checklist. If a thought was checked, it was scored as 1; otherwise it was scored as 0.

Representation of Output Features

Each element of each Broken Date message was classified as belonging to one of 72 message-frames. Based on these codings, each message was represented as a vector of binary values on the 72 variables constituted by the possible frames. If a particular message contained a particular frame, the value of the bit representing that frame was set to 1; otherwise, it was set to 0. When a frame occurred more than once in a single message, duplicate occurrences were disregarded.[33]

Application of these procedures created a data set containing, for each subject (N = 320), a 121-bit input vector representing the level of activation of each thought from the thought checklist and a 72-bit output vector representing the level of activation of each message-frame.

---

[31] Message-frames are not in people's heads. They are an artifact of the analysis used to group synonymous message elements produced by different individuals. A person only has message elements in his/her head, not message frames, and therefore people do not have to instantiate abstract types.

[32] Such models are often called "supervised" because learning relies on input from a "teacher" regarding correct examples. These models are contrasted with unsupervised techniques that need no teacher.

[33] Duplicate frames are disregarded on theoretical grounds. It is argued that people have a set of active frames and that they simply read off the active frames at any given moment. If certain frames remain active, they may be expressed more than once.

The Learning Algorithm

The proposed architecture for this simulation is a standard, feedforward network trained by a fast version of the back-propagation learning rule. The following sections describe some important details of the model.[34]

Activation rule. The net input to a given unit *i* is simply a sum of the products of incoming weights and activations for a given pattern *p*:

(15) $\text{net}_{pi} = \bullet w_{ij} a_{pj}$

where w_{ij} is the weight connecting unit j to unit i and a_{pj} is the activation of unit *j* at pattern *p*. The activation rule for each unit in the model is given by the so-called logistic function:

(16) $a_{pi} = 1/(1 + e^{-\text{net}pi})$

where a_{pi} is the activation of unit *i* upon presentation of pattern *p*, and net_{pi} is the net input to unit *i* for pattern *p*. This function yields a smooth sigmoidal curve ranging from 0 to 1 (exclusive).

Back-propagation learning rule. The general form of learning in supervised pattern associators involves three steps for each pattern presented: (a) feedforward passing of activation from inputs (possibly to hiddens) to outputs; (b) calculation of the difference between target and actual output values (error); and (c) weight modification by back-propagation of error through the network. The general form of the weight modification rule for back-propagation (and several other supervised pattern associators) is given in (17):

(17) $w_{ij} = \text{eta}(\text{delta}_{pi} a_{pj})$

where eta is a parameter controlling the size of each learning and delta_{pi} is a measure of the difference between the target and actual activations of unit *i* (explained below). When hidden units are present, separate delta values must be calculated for output units and for hidden units, since information about the target activation of hidden units is not directly available.[35] Thus, delta_{pi} for any output unit is given in (18):

(18) $\text{delta}_{pi} = (t_{pi} - a_{pi})a_{pi}(1 - a_{pi})$

where t_{pi} is the target activation of some output unit *i* and a_{pi} is the actual activation produced by that unit for pattern *p*. Delta_{pi} for any hidden unit is given in (19):

(19) $\text{delta}_{pi} = a_{pi}(1 - a_{pi})\bullet\text{delta}_{pk} w_{ki}$

where delta_{pk} is the delta for some output unit *k* at pattern *p*, and w_{ki} is the weight from hidden unit *i* to output unit *k*. Since there are no explicit targets for hidden units, the error signal for each hidden unit is a function of the weighted sum of all the errors it contributed to at the output layer.

---

[34] The mathematical details of the model are adapted from McClelland & Rumelhart (1989, esp. ch. 5) and Rumelhart, Hinton, & Williams (1986).

[35] A hidden unit is a type of unit that intervenes between input and output units. Its function is analogous to an interaction term in polynomial regression. The presence of hidden units greatly increases the representational power of a network, making it possible to associate patterns that are related to one another in nonlinear ways.

Speeding Up Back-Propagation: The Quickprop Algorithm

Back-propagation has proved useful on a wide range of problems, but it is a very slow algorithm. Fahlman's (1988) Quickprop algorithm was designed to remedy the learning speed problems inherent in standard back-propagation. This section describes the Quickprop algorithm, a method that has been shown to speed up standard back-propagation by an order of magnitude or more (Fahlman, 1988).

The flat spot. The first problem Fahlman noticed with standard back-propagation pertains to the calculation of the error signal for each unit. Note that the equations in (18) and (19) both contain the quantity $a_{pi}(1 - a_{pi})$. This quantity is the first derivative of the sigmoidal activation unction, sometimes referred to as sigmoid-prime (Fahlman, 1988). Sigmoid-prime approaches zero as the activation of unit *i* at pattern *p* approaches either zero or one. Consequently, the error signal passed back through the network will be small whenever the unit in question has a value approaching either zero or one. Consider the case where the actual activation of a unit is 0.99 and it should be 0.00. Due to the nature of sigmoid-prime, a small error signal will be sent when a large error signal should be sent. Since the change in weight to a given unit is proportional to the error signal from that unit, small error signals translate into small weight changes.

The upshot is that the "flat spot" in sigmoid-prime may cause units far from their desired values to get stuck for long periods during training, needlessly slowing convergence. Fahlman's (1988) solution is simple. Quickprop adds a small constant quantity (normally 0.1) to the sigmoid prime function and thereby keeps units from getting stuck in the flat spot.

The quadratic estimate. The essence of Quickprop, however, consists not in removing the flat spot but in fundamentally altering the method by which the algorithm converges to a solution. Back-propagation is one of a class of algorithms that reaches its solution by (a) estimating the effect of a change in a given weight-value on the error of a given unit (i.e., the local slope of the error surface with respect to a given weight) and (b) changing the weight to decrease the error. This is the method of gradient descent, where the gradient is on the error surface plotted in weight-space. Typical gradient descent methods take steps down the gradient that are proportional to a learning rate parameter typically called eta (see (10)). The frequently small values of eta needed to descend the error gradient successfully impede fast learning in standard back-propagation. According to Fahlman (1988, p. 10), "we do not want to take infinitesimal steps; we want to take the largest steps possible without overshooting the solution." Therefore, Quickprop assumes that, for a given weight, the error surface with respect to that weight is well approximated by a parabola. Making this assumption, the function for changing a given weight becomes a simple function of the current and previous slopes of the error surface as well as the previous weight change:

$$(20)\quad w(t) = [S(t) / (S(t-1) - S(t))]\, w(t-1)$$

where w(t-1) is the previous change in weight, S(t) is the current slope of the error surface, and S(t-1) is the previous slope of the error surface.[36]

When the current slope is less than the previous slope and in the same direction, the weight is changed in the same direction in proportion to the error reduction achieved by the previous step. When the current and previous slopes have opposite signs, the weight change is in the opposite direction of the previous step. When the current slope is greater than the previous slope, one is ascending rather than descending the error surface. In this case, a fraction of the previous step is used as the current step, and the quadratic estimate is not used.

[36] Subscripts are dropped here for readability.

To start the process or to calculate the step when the slope is unchanged, standard gradient descent is used (see Fahlman, 1988).

Architecture of the Model

The architecture of a network model refers to the number, type, and connectivity pattern of units in a network. The best architecture for a given learning problem is determined empirically. For this analysis, several architectures were evaluated before settling on a given model to analyze. This section describes the methods used for evaluating alternative architectures and presents performance data for the architecture that was eventually selected (and for the first runner-up).

The purpose of this project is to understand the relationship between what people are thinking and what they say. The purpose of training a machine learning model to approximate this mapping is not just to have a black box that takes thoughts as input and outputs message-frames. Rather, we want to find out what is learned. In a PDP model this amounts to analyzing the "knowledge" in the network. Knowledge is embodied in a matrix of learned weights connecting thoughts to message-frames. So training the model is just a prelude to analyzing the learned weights.

Training and testing. Two criteria are used to evaluate whether the weights learned by a given model are worth examining: (a) performance on the entire example set, and (b) average performance on cross-validation example sets.

To be confident in the generality of the conclusions drawn from an analysis of the learned weight matrix, one must demonstrate that the network performs well on cases that it has not been trained on. That is, in order to be confident in the generality of one's analysis, there must be evidence that the network has extracted important underlying relationships between thoughts and messages, not just idiosyncratic relations embodied by a given example set. This is particularly important since some PDP architectures have the capacity (in effect) to memorize their training data. In such a case, performance on the training examples is perfect, but generalization to novel cases is poor. In this investigation, several network architectures were evaluated. Many were able to learn the entire data set perfectly. Therefore, the task was to choose from among these the architecture with the best generalization performance.

Definition of performance measures. Before proceeding some definitions of performance measures are in order. An output vector consists of 72 elements, each of which has a value of either 0 or 1. A bit (or element) is defined to be in error if the difference between the desired and actual value of the bit exceeds a fixed threshold. In this investigation, the threshold is set to .4. Bits that are supposed to be 0 are counted as correct if their actual value is less than or equal to .4. Bits that are supposed to be 1 are counted as correct if their actual value is greater than or equal to .6. In all other cases, a bit is in error.[37] An epoch is defined as one pass through all of the examples in a given set. Total error bits refers to the total number of error bits summed across all the patterns in an epoch. The criteria for selecting the network architecture can now be stated. The network to be chosen will produce 0 total error bits when trained on the entire example set and will have better generalization performance than any other architecture.

Network Performance

Search through the space of possible architectures began with the simplest possible architecture, a network with direct connections from input to output with no hidden units. Such a network was not able to learn the entire example set perfectly. Next, a network with direct connections from input to output and a single hidden unit was evaluated. This architecture was able to learn the entire example set perfectly. The entire example set was

[37] These are the values recommended by Fahlman (1988).

also perfectly learned by a network with 23 hidden units and no direct connections between input and output.[38] The generalization performance of these two networks was evaluated.[39]

Cross-validation procedure. To evaluate the generalization performance of each candidate network, a form of hold-out procedure was used (see, for example, Lachenbruch & Mickey, 1968). In the most extreme case N-fold cross-validation (i.e., hold one out) involves training a classifier on N-1 cases, testing on 1 case, and repeating this N times. The procedure yields a more accurate estimate of true classification performance than (say) training on half the data and testing on the other half. Conducting 320 experiments for each generalization result was not feasible, so 10-fold cross-validation was performed.

The examples were partitioned into ten different subsamples. In each cross-validation experiment, 90% of the examples ($n = 288$) were used to train the network, and 10% ($n = 32$) were used to test it. Naturally, a different 10% was chosen as the test set for each experiment. In all, each example was part of the training set in 9 experiments and part of the test set in one. The generalization performance of a given model was given by the percentage of error bits averaged across the ten cross-validation experiments.[40]

Performance results. Performance data for the best two architectures are given in Table 6 and Table 7. The network with one hidden unit and direct connections from inputs to outputs yielded the best overall generalization performance. Across all ten cross-validation experiments, this network correctly classified an average of 87.8% of the output features, significantly greater than the 86.4% achieved by the network with 23 hidden units and no direct connections between input and output ($t(9) = 4.469$, $p < .002$, $r = .83$). This high level of generalization performance must be qualified, however.

Recall from chapter six that the average number of message-elements per subject was 5.44. In other words, the average output vector had approximately 5 bits turned on and 67 turned off. In this circumstance, the network learns that most output features are off most of the time. Therefore, we should expect that the network will perform fairly well on novel cases, because they too will have these "sparse" output patterns (i.e., many off, few on).

A stricter measure of generalization performance asks what percentage of the units that were supposed to be on were actually turned on. We will refer to this measure as the on-hit accuracy. The same architecture that had an average overall accuracy of 87.8% had an average on-hit accuracy of only 22.1%. This is disappointingly low. For novel thought patterns, only slightly more than one out of five message-frames that are supposed to be on are actually turned on by the network. The small number of examples in the current sample undoubtedly accounts for part of this poor generalization.

---

[38] Networks with more hidden units were also able to learn the entire set perfectly. The network with the fewest hidden units was chosen as a basis for comparison because it is generally recognized that smaller networks yield better generalization performance, the idea being to fit the data using the fewest degress of freedom.

Networks with fewer than 23 hidden units and no direct connections between input and output could not learn the entire data set perfectly and were therefore excluded.

[39] There are several adjustable parameters in Quickprop. The following parameters deviated from Fahlman's (1988) defaults: for the model with 1 hidden unit, epsilon = 0.05, MaxFactor = 1.75. For 23 hidden unit model, epsilon = 0.1, MaxFactor = 0.5.

[40] Percentage error bits is a ratio of total error bits to total bits possible, where total bits possible is the product of the number of bits per example and the number of test examples. In this case, there are 72 bits per output vector and 32 test cases (total bits possible = 2304)

Table 6

Performance Measures for Best Network Architectures

	% On-Hits		% Total Hits	
	No. Hidden Units		No. Hidden Units	
Exp. No	1	23	1	23
1	15.0	18.5	88.1	87.3
2	25.0	20.2	88.2	87.0
3	23.0	21.9	87.6	83.8
4	19.1	27.4	87.4	86.9
5	25.0	16.2	87.8	86.6
6	27.1	18.8	88.8	86.6
7	23.4	22.7	88.0	87.3
8	22.1	20.4	88.5	86.3
9	22.2	17.8	86.6	85.5
10	19.2	16.7	87.1	86.4
Mean	22.11	20.1	87.8^{a}	86.4^{b}
S.D.	3.53	3.33	0.66	1.05

Note. Means with different subscripts differ significantly at $p < .005$.

Relationship between generalization performance and absolute frequency of frames. Rather than thinking of the network's task as predicting one 72-bit vector for each thought pattern, it can be thought of as 72 independent prediction problems: given a thought pattern, is each message-frame on or off?[41] Viewed in this light, the network understandably has difficulty generalizing about low frequency message-frames. There simply are not enough positive examples to learn from. To illustrate the problem, it is worth noting that more than half of the 72 message-frames occurred less than 20 times (see Figure 6) in the entire sample.

In order to get a quantitative grasp on this relationship, it is useful to examine the relationship between per-frame generalization performance and frame frequency. According to the analysis just offered, we expect strong relationships between on-hit accuracy and frame frequency, with a better accuracy on more frequent frames. To examine this relationship, generalization performance was broken down into three scores: (a) overall accuracy, (b) on-hit accuracy, and (c) off-hit accuracy. Overall accuracy is just a ratio of the number of bits correctly classified to the total number of bits attempted. On-hit accuracy refers to the

[41] Scott Fahlman (personal communication) pointed this out.

proportion of bits that were on when they were supposed to be on. Off-hit accuracy refers to
Table 7

Number of Epochs to Learn Training Set Perfectly

	No. Epochs	
	No. Hidden Units	
Experiment No.	1	23
1	6075	3745
2	1320	3215
3	1550	3830
4	1440	4265
5	1030	4315
6	2970	3420
7	1855	2910
8	1525	3220
9	1410	4075
10	3930	3525
Mean	2310.5	3688
S.D.	1592.39	536.52

the proportion of bits that were off when they were supposed to be off.

Performance was broken down on a frame-by-frame basis. As expected, there was a strong positive relationship between on-hit accuracy and frame frequency ($r(44) = .567$, $p < .001$). On-hit accuracy was better for more frequent frames. In contrast, there was a very strong negative relationship between off-hit accuracy and frame frequency ($r(63) = -.933$, $p < .001$) and between overall accuracy and frame frequency ($r(70) = -.975$, $p < .001$).[42]

If the network learns that most frames are off most of the time, these relationships are expected. Low-frequency frames are accurately predicted because the network predicts they will always be off. High-frequency frames are less accurately predicted because they contradict the trend that the network learns. This does not mean that overall accuracy will continue to decrease as the number of examples (and hence the absolute frequency of each

[42] On-hit raw proportions underwent a natural logarithmic transformation before being entered into the analysis (to restore normality). Off-hit and total-accuracy proportions underwent a hyperbolic arctangent transformation before being analyzed. Degrees of freedom differ because some frames had zero on-hits or off-hits and therefore could not be transformed.

frame) increases. The relationship between overall accuracy and frame frequency is probably U-shaped, with high accuracy at low and high frequencies (see Figures 7 and 8). To sum up, we should expect overall performance to improve with a larger example set. This is a truism in machine learning.

It would be preferable to show better generalization results, but the low percentage of on-hits does not severely undermine the goal of this chapter. The goal is to establish the face validity of our modelling strategy and to evaluate its potential, not to build the beat-all-and-end-all model. It simply did not make sense to gather thousands of examples before the basic usefulness of the new methods (of thought elicitation and message-element coding) had been established. Nevertheless, any conclusions drawn from the analysis of learned weights presented below must be qualified by the relatively low level of on-hit generalization.

Example messages output by network. To give the reader a feel for the network's generalization performance before proceeding to analyze what the network learned, some example messages generated by the network are displayed. These messages were generated in response to thought-patterns that were not part of the training set.[43] Thus, they do not perfectly predict or mirror the target messages.

Each example message was generated by a three-step procedure. First, a novel thought pattern was presented to the trained network and activation was allowed to pass forward to the output units. Next, message frames corresponding to output units with levels of activation exceeding a given threshold (0.6) were selected. Finally, for each message frame selected, an instance of that frame was randomly selected from the sorted message-elements (see chapter 6) and output. Displayed in Table 8 are the following: (a) the original message produced by study participants in connection with a given pattern of reported thoughts, (b) the message-elements produced by randomly instantiating each message-frame, (c) the set of message-frames corresponding to the original message, and (d) the set of message-frames produced by the PDP model.[44]

Analysis of Weights

Methods

We can conclude from the relatively weak generalization performance that the underlying relationships that connect thoughts to messages were not fully extracted by the current model. Therefore, an exhaustive analysis of the learned weight matrix is put off until a model can be trained with a largerexample set. It is still possible, however, to assess the face validity of the model by a rough examination of the projective and receptive fields of the input and output units respectively.[45] The next section describes a method for analyzing projective and receptive fields and carries out such an analysis for selected input and output units. The weights analyzed below come from a network with one hidden unit and direct connections from input to output. The network was trained to perfection on the entire example set (N = 320). Thus the weights analyzed are not from a cross-validation model but rather from a model based on all available data.

Receptive fields. In neurobiology, the receptive field refers to a range of sensory stimuli or "region of the receptive surface" (e.g., an area on the skin; Kandel & Schwartz,

[43] The weight matrix from the sixth cross-validation experiment was used.

[44] To repeat, the random instantiation of message-frames as concrete message elements is not a process that occurs in people. People have concrete message elements in mind, not message-frames. This is just a way of illustrating how a randomly selected person might respond to a given pattern of thoughts.

[45] This mode of analysis was suggested to me by Scott Fahlman. Use of the terms "receptive and projective fields" is based on an analogy to sensory neurophysiology.

1985, p. 288) to which a given neuron is highly sensitive. Here, a given output unit's excitatory receptive field is comprised of the 12 input units (ten percent of the total number

Table 8

Comparison of Predicted Versus Actual Message Frames and Model-Generated Versus Actual Message Instances

---

Actual message (#166)

Maybe I am just not your friend.
If you do not have time for friends maybe you are too busy.
You have blown me off every time we have tried to go out these past weeks.
I think you have no idea what the word commitment means.
Making a date for tomorrow night might be too difficult a commitment for you.

Model-generated message

I am really hurt about what you have been doing.
What is going on in your life that makes you so scatterbrained that you can not even plan to meet with me without eventually cancelling.
Now you are just not there for me.
I will go by myself.
Something always comes up that is more important that me to you.
Are you going to cancel out on me again.
Maybe I will be able to see you in the future.

Frames for actual message

You have broken our dates a lot lately.
Your behavior is bothering me.
You've been busy.
You have not been a good friend lately.
I expect you to cancel tomorrow, too.

Frames generated by model

Your behavior is bothering me.*
What is going on?
You have not been a good friend lately.*
I'll find alternatives.
You have made me a low priority.
I expect you to cancel tomorrow, too.*
Check with me another time.

Actual message (#167)

Call me when you are free sometime.

I miss doing things with you.
I hope we can plan something to do sometime.

Table 8 (cont'd.)

I understand that you are a very busy person.
It is all right.

Model-generated message

I still consider you a very good friend.
When you decide you have some time call me.
I know people.
Could I go to the meeting with you tonight.

Frames for actual message

You've been busy.
I want to get together.
I miss you.
Call me when you have more time.
Okay, I can go tomorrow.

Frames generated by model

We have been good friends for some time.
Call me when you have more time.*
I understand your constraints.
Let's find a compromise plan.

Actual message (#168)

We will see what happens.

Model-generated message

Terry I am tired of having you cancel out all the time.
Please do not continue treating me as you have.
Try to realize that I have a life to.
I have to do some other things.
If you cancel out tomorrow so help me I will kill you.
I will go to the movie tomorrow night if you know absolutely that you will be there.

Frames for actual message

Tomorrow is fine if we can make it.

Frames generated by model

I'm tired of your behavior.

Table 8 (cont'd.)

Your behavior is bothering me.
I have a busy life, too.
Sorry, I can't go tomorrow.
If you cancel again, I'll be angry.
Tomorrow is fine if we can make it.*

Actual message (#169)

Every time I am looking forward to something you cancel out.
It would be nice if you would keep a few of our engagements.
That is okay terry.

Model-generated message

When you cancel plans that sometimes means I miss out on other things I could have dome.
You will not hear from me in the meantime.
If you have a problem going out with me then stop making plans with me.
I have plans tomorrow.
Maybe I will have time to go.

Frames for actual message

You have broken our dates a lot lately.
Keep our date tomorrow.
Okay, I can go tomorrow.

Frames generated by model

I've made sacrifices to be with you.
Since you cancel, I won't make plans with you.
If you don't want dates don't make dates.
Sorry, I can't go tomorrow.
Check with me another time.

Actual message (#170)

I hope you can find time in your busy schedule to come with me.
All right.

Model-generated message

Somehow our plans always get cancelled.
I hate not to take advantage of my free time.

Table 8 (cont'd.)

I find it a little fishy that every time I have a date with you you seem to have something come up.
Do not make plans when you are more than likely to back out.
If you want to go out tomorrow call me.

Frames for actual message

I hope we get together tomorrow.
Okay, I can go tomorrow.

Frames generated by model

You have broken our dates a lot lately.
Your behavior is bothering me.
I suspect that you are avoiding me.
If you don't want dates don't make dates.
Call me when you have more time.

Note. An asterisk indicates a correctly predicted frame.

of inputs) with the strongest positive weights to that output.[46] Looking at the receptive fields of selected message-frames, we can see which thoughts strongly excite or inhibit them. In this manner, we can get a partial picture of the functionality of the network. Receptive fields of the 5 message-frames with the highest on-hit generalization performance are analyzed.

Analysis of Frame 54

The message-frame with the best generalization performance was frame 54: "Sorry, I can't go tomorrow". Table 9 displays the inhibitory and excitatory receptive fields for this frame. Each element of a receptive field is defined by a thought and by the magnitude of the weight projecting from that thought. For example, the thought "Terry would respect me for being firm" had the most influence on this frame, with a positive weight of 19.41.

Excitatory receptive field. Frame 54 is the message-frame corresponding to a refusal of Terry's request to reschedule the date. On a standard plan-based account, a refusal would be triggered by perceived violations of the felicity conditions on requesting. A refusal would be prompted if, for example, a request were deemed insincere, or the requested party were not able to grant the request, etc. Alternatively, a sincere request could be refused if the hearer is not willing or able to do the requested act.

In the excitatory receptive field for frame 54, several of the conventional grounds for refusing a request are present. "Terry's using activities as an excuse not to see me" is a thought questioning Terry's sincerity. Both "We shouldn't make definite plans again until Terry has more free time" and "Terry is disorganized" pertain to Terry's ability to do his part of the requested act. As for the hearer, "I should stop seeing

[46] Inhibitory fields are defined similarly for the strongest negative weights.

Terry for a while" and "I should refuse to go with Terry tomorrow" concern the hearer's intention and/or desire to do the requested act. The presence of each of these thoughts is correctly predicted by a standard, speech-act based model.

However, there are thoughts that strongly excite frame 54 that are not predicted by a plan-based theory. The presence of "Terry's guilty" and "Terry's rude" suggest that some people refuse the request simply to punish Terry for past bad behavior. Others refuse the request because by refusing they convey a certain impression of themselves to Terry (e.g.,

Table 9

Excitatory and Inhibitory Receptive Fields for Frame 54

"Sorry, I can't go tomorrow"
(Total-Accuracy = 69.37%, On-Hit Accuracy = 56.8%)

Excitatory Receptive Field

Weight	Thought
19.414	Terry would respect me for being firm
19.281	Terry's guilty
17.355	Terry's using activities as an excuse not to see me
11.951	I should stop seeing Terry for a while
10.664	I should refuse to go with Terry tomorrow
10.370	Terry ought to take responsibility for our friendship
10.024	Terry feels bad about cancelling the date
9.743	Terry acts responsibly about attending meetings
9.613	We shouldn't make definite plans again until Terry has more free time
8.068	Terry is rude
7.994	We should get together at a time of Terry's choosing
7.571	Terry is disorganized

Inhibitory Receptive Field

Weight	Thought
-41.307	Terry cares about me and my feelings
-16.480	Hidden unit
-15.931	Terry is insecure
-15.354	Terry would like it if the situation were cleared up
-14.977	Terry is highly motivated
-14.681	Terry is an overachiever
-12.443	Terry was considerate to call me
-11.637	Terry has personal problems and obligations I don't know about
-10.997	Terry and I could do something else and skip the movie
-10.992	Terry will try to keep the appointment tomorrow night
-10.706	Terry is insensitive to the feelings of others
-10.321	If I criticize Terry it will hurt our friendship

"Terry would respect me for being firm"). Or they use the refusal to redefine the structure of

rights and obligations in the relationship (e.g., "Terry ought to take responsibility for our friendship," and "We should get together at a time of Terry's choosing").

Inhibitory receptive field. As expected, frame 54 is strongly inhibited by thoughts affirming Terry's sincerity (e.g., "Terry cares about me and my feelings," "Terry would like it if the situation were cleared up," "Terry was considerate to call me," "Terry will try to keep the appointment tomorrow night"). Some thoughts inhibit frame 54 due to its potential negative impact either on Terry's self-image or on the friendship (e.g., "Terry is insecure," "If I criticize Terry it will hurt our friendship"). Finally, frame 54 is inhibited by thoughts acknowledging Terry's unobjectionable (prosocial) motivations (e.g., "Terry is highly motivated," "Terry is an overachiever") and by thoughts attributing Terry's actions to unknown causes (e.g., "Terry has personal problems and motivations I don't know about").

Analysis of Frame 49

Excitatory receptive field. Table 10 displays the accuracy data and receptive fields for frame 49, "Call me when you have more time." This frame is excited by obviously relevant thoughts about future plans with Terry (e.g., "Terry should be the one to make plans in the future") one of which is a paraphrase of the frame itself ("Terry should call me back when he/she has time for a long talk"). Negative characterizations of Terry excite this frame (e.g., "Terry is rude," "Terry knows too many people"), but so do certain positive characterizations (e.g., "Terry feels bad about cancelling the date," "Terry is highly motivated," "Terry thinks it is important to keep commitments"). Thoughts expressing frustration are excitatory (e.g., "I'm tired of being patient with Terry," "It would be better if Terry weren't so sure of my friendship"), indicating that for some speakers this frame signals frustration. In contrast, thoughts expressing concern about the impact of frame 49 are also excitatory (e.g., "If I criticize Terry it will hurt our friendship," "Terry is insecure"). That the same frame is excited by seemingly contradictory thoughts supports the contention that the expression of a given content element may have very different functional significance for different people.

Inhibitory receptive field. Thoughts about Terry's sincerity (e.g., "Terry will try to keep the appointment tomorrow night") and prosocial motivations (e.g., "Terry is an overachiever," "Terry acts responsibly about attending meetings") inhibit frame 49, suggesting that some people view the frame as an inappropriate "kiss off" in light of Terry's good intentions. Similarly, thoughts expressing understanding inhibit this frame (e.g., "I can make it tomorrow," "Terry should know that I understand about the cancellations," "I should be understanding about Terry's problems"). Two negative thoughts about Terry (e.g., "Terry should be put on the defensive," "Terry is insensitive to the feelings of others") are also inhibitory. To some message producers telling Terry to call when s/he has more time is apparently not an effective way to put Terry on the defensive or punish Terry for being

Analysis of Frame 65

Excitatory receptive field. As expected, frame 65, "I want to see the movie tonight" is excited by the nearly identical thought "I really want to see this movie" (see Table 11). Excitation also comes from a set of negative thoughts about Terry (e.g., "Terry's a liar," "Terry is inconsiderate," "Terry is guilty," "Terry irritates me," "Terry should apologize for being a bad friend"), as well as from a thought about the impression created by expressing frame 65 ("Terry would respect me for being firm"). This pattern of excitation suggests that people who have ill will toward Terry are likely to express their own wants. In this context, such an expression may also function as an indirect refusal of the request.

Inhibitory receptive field. The inhibitory receptive field of frame 65 is difficult to interpret. Apparently, the simple expression of a want is not thought to be helpful in getting Terry to behave because thoughts about what Terry should do inhibit this frame (e.g., "Terry should put more effort into our friendship," "Terry should promise not to cancel again,"

"Terry ought to take responsibility for our friendship"). Neither is this frame useful for signalling anger and frustration (e.g., "Terry needs to know I'm a little angry," "Terry doesn't respect me," "Terry's irresponsible"). Some message producers seem to

Table 10

Excitatory and Inhibitory Receptive Fields for Frame 49

"Call me when you have more time"
(Total-Accuracy = 70.00%, On-Hit Accuracy = 48.65%)

Excitatory Receptive Field

Weight	Thought
19.264	Terry is rude
18.511	Terry should be the one to make plans in the future
17.953	Terry knows too many people
15.108	Terry feels bad about cancelling the date
13.565	Terry should skip the meeting tonight
12.237	Terry is highly motivated
11.955	Terry should call me back when he/she has time for a long talk
11.078	Terry thinks it is important to keep commitments
10.798	I'm tired of being patient with Terry
10.654	If I criticize Terry it will hurt our friendship
10.285	It would be better if Terry weren't so sure of my friendship
10.011	Terry is insecure

Inhibitory Receptive Field

Weight	Thought
-30.548	Bias unit[a]
-22.978	Terry acts responsibly about attending meetings
-19.113	Terry should be put on the defensive
-18.529	I can make it tomorrow
-16.485	Terry is an overachiever
-16.139	Talking with Terry about the friendship is important
-16.057	If Terry knows I feel hurt, it will make me look bad
-14.012	Terry will try to keep the appointment tomorrow night
-11.844	Terry would like it if the situation were cleared up
-11.826	Terry is insensitive to the feelings of others
-11.236	Terry should know that I understand about the cancellation
-11.015	I should be understanding about Terry's problems

---

[a]The bias unit is analogous to the constant term in a regression equation. In a PDP model, the bias is a unit with a constant activation of 1. The weights from the bias unit to each hidden and output unit are modified during training as are all other weights.

view this element as damaging since it is inhibited by positive thoughts about preserving or repairing the friendship (e.g., "We need to get things out in the open," "Terry wants to do the movie tonight but can't because of a meeting," "If I criticize Terry it will hurt our friendship," "I want to spend time with Terry").

Table 11

Excitatory and Inhibitory Receptive Fields for Frame 65

"I want to see the movie tonight"
(Total-Accuracy = 90.31, On-Hit-Accuracy = 39.13)

Excitatory Receptive Field

Weight	Thought
8.751	Terry would respect me for being firm
7.748	Terry's a liar
7.677	I really want to see this movie
7.114	Terry is inconsiderate
6.884	Terry thinks it is important to keep commitments
6.004	I should let Terry know how I feel about the cancellation
5.981	Terry's guilty
5.361	Terry has personal problems and obligations I don't know about
4.954	Terry irritates me
4.716	Our friendship needs to be put back on track
4.651	If I refused to see Terry, Terry would miss me
4.158	Terry should apologize for being a bad friend

Inhibitory Receptive Field

Weight	Thought
-13.503	Bias unit
-8.755	Terry doesn't respect me
-8.046	Terry needs to know that I am a little angry
-7.950	We need to get things out in the open
-7.058	Terry should put more effort into our friendship
-6.161	Terry's irresponsible
-6.118	Terry should promise not to cancel again
-5.632	I'm tired of being patient with Terry
-5.407	Terry ought to take responsibility for our friendship
-5.377	Terry wants to do the movie tonight but can't because of a meeting
-5.372	If I criticize Terry it will hurt our friendship
-5.212	I want to spend time with Terry

Analysis of Frame 6

Excitatory receptive field. The receptive fields of frame 6, "Is there a problem?," are displayed in Table 12. Just as one would expect, this frame is strongly excited by thoughts that anticipate or identify a problem (e.g., "Terry has a problem with me that he/she is afraid to bring up," "Terry wants to do the movie tonight but can't because of a movie," "If Terry has a problem with me, he/she should say so," "Terry wants to break off our friendship"). This pattern suggests that frame 6 functions straightforwardly as a question for some people. Others may think that asking for an account makes them look tough (e.g., "Terry would respect me for being firm"). Positive thought about Terry and his sincerity suggest that still others may intend this frame to signal concern for Terry (e.g., "I should be supportive to Terry," "Terry is a good friend," "Terry will try to keep the appointment tomorrow night"), while negative thoughts exciting this frame imply that asking for an account may be a form of punishment (e.g., "Terry doesn't consider the needs/feelings of others to be important," "I want to do to Terry what he/she did to me").

Inhibitory receptive field. For some people, the expression of frame 6 has damaging consequences that might prevent or postpone future meetings. For example, thoughts about showing understanding (e.g., "Terry wants me to understand what is happening in this situation") and about avoiding criticism (e.g., "If I criticize Terry it will hurt our friendship") inhibit the request for an account, as do thoughts about seeing Terry soon (e.g., "I want to spend time with Terry," "Terry would like to see me tomorrow night"). Also, when the message producer has clear ideas about what Terry should do (e.g., "Terry should skip the meeting tonight," "Terry should put more effort into our friendship") this frame is inhibited.

Analysis of Frame 1

Excitatory receptive field. Table 13 displays the receptive fields and generalization accuracy data for frame 1, "You have broken our dates a lot lately." The presence of certain thoughts about Terry's motives (e.g., "Terry is highly motivated," "Terry is an overachiever," "Terry cares about me and my feelings," "Terry will try to keep the appointment tomorrow night") and about proper friendship behavior (e.g., "I should be understanding about Terry's problems," "Talking to Terry about the friendship is important") suggests that frame 1 may function as a friendly reminder to Terry, who in a frenzy of activity, has failed to realize what s/he has been doing to his/her friend. The other excitatory thoughts reflect anger and resentment (e.g., "I want Terry to feel guilty," "Terry's irresponsible," "Terry needs to know that I'm a little angry," "Terry's actions hurt my feelings") and imply that frame 1 may also function as an
accusation.

Inhibitory receptive field. Frame 1 may signal bad feelings for some because it is inhibited by thoughts about being understanding (e.g., "Terry expects me to be understanding") and by thoughts about the need to hide bad feelings (e.g., "If Terry knows I feel hurt it will make me look bad," "I don't want Terry to know I feel hurt"). It is difficult to interpret the remaining inhibitory thoughts.

Goals as Distributed Patterns of Thoughts

It seems quite helpful when interpreting the receptive fields to identify coalitions of excitatory and inhibitory thoughts that share conceptual characteristics. Thus the analysis above includes references to conceptually grouped clusters of thoughts (i.e., positive thoughts about Terry, thoughts about proper friendship behavior, thoughts about Terry's motives, etc.). However, just because subsets of thoughts share conceptual features does not mean that thoughts actually co-occur in neat conceptual groups. In order to substantiate the analysis of excitatory and inhibitory receptive fields and to examine the actual patterns of co-

occurrence among thoughts, a cluster analysis was performed on the 121 variables from the thought-checklist.

Table 12

Excitatory and Inhibitory Receptive Fields for Frame 6

"Is there a problem?"
(Total-Accuracy = 86.25, On-Hit Accuracy = 35.48)

Excitatory Receptive Field

Weight	Thought
15.127	Terry has a problem with me that he/she is afraid to bring up
11.377	Terry wants to do the movie tonight but can't because of a movie
10.709	Terry doesn't consider the needs/feelings of others to be important
9.778	Terry should keep appointments
9.518	I should be supportive to Terry
9.328	Terry tries to be too involved
8.368	I should be firm in dealing with Terry
7.609	If Terry has a problem with me, he/she should say so
7.404	I want to do to Terry what he/she did to me
6.279	Terry is a good friend
5.729	Terry will try to keep the appointment tomorrow night
5.651	Terry wants to break off our friendship

Inhibitory Receptive Field

Weight	Thought
-16.709	Bias unit
-10.720	Terry is self-centered
-10.085	Terry wants me to understand what is happening in this situation
-9.576	Terry should skip the meeting tonight
-7.812	Terry knows too many people
-7.770	If I criticize Terry it will hurt our friendship
-7.690	I want to spend time with Terry
-7.048	Terry's avoiding me
-6.952	We need to get things out in the open
-6.316	Terry should put more effort into our friendship
-6.193	Hidden unit
-6.183	Terry would like to see me tomorrow night

Clustering Method

Cluster analysis is a notoriously flexible procedure for examining groupings in data. Similarity-based clustering techniques are especially sensitive to the manner in which similarity is defined. In fact, the clusters discovered by similarity-based techniques may be as much a reflection of the chosen similarity metric as of any "natural" grouping in the data. In an effort to avoid introducing an unnecessary arbitrary element into the present analysis, the clustering technique used here is based on a principal component analysis of the correlations between variables.[47]

Table 13

Excitatory and Inhibitory Receptive Fields for Frame 1

"You have broken our dates a lot lately"
(Total-Accuracy = 59.37%, On-Hit Accuracy = 34.02%)

Excitatory Receptive Field

Weight	Thought
24.979	Hidden unit
13.485	Terry is highly motivated
9.163	Terry cares about me and my feelings
7.812	I want Terry to feel guilty
6.637	Terry is an overachiever
6.490	I should be understanding about Terry's problems
6.160	Talking with Terry about the friendship is important
5.544	I want Terry to go out with me
5.182	Terry will try to keep the appointment tomorrow night
5.179	Terry's irresponsible
4.926	Terry needs to know that I am a little angry
4.627	Terry's actions hurt my feelings

Inhibitory Receptive Field

Weight	Thought
-23.403	Bias unit
-10.734	Terry's guilty
-7.978	Terry expects me to be understanding
-7.348	Terry feels bad about cancelling the date
-6.279	Terry is selfish
-5.399	Terry's telling the truth
-5.302	Terry should call me back when he/she has time for a long talk
-4.890	Terry should be put on the defensive
-4.532	If Terry knows I feel hurt, it will make me look bad
-3.856	Terry doesn't realize that cancelling appointments can hurt others
-3.844	Terry should put more effort into our friendship

[47] It might be noted that correlation is just another similarity metric, and this is indeed true. However, it is the standard measure of similarity upon which linear statistical models are based.

-3.768 I don't want Terry to know that I feel hurt

Data for the cluster analysis were 320 observations on 121 variables (thoughts from the thought-checklist task). Disjoint clusters were computed using SAS's VARCLUS procedure, "a type of oblique component analysis related to multiple group factor analysis" (SAS Institute, Inc., 1985, p. 802; see also Harmon, 1976). Roughly speaking, VARCLUS works in six steps: (a) compute the principal components of the correlation matrix for all variables, (b) use the first and second principal components as the variates for two new clusters, (c) assign each variable from the original set to the cluster variate with which it has the highest squared correlation, (d) recalculate principal components for all clusters, (e) create a new cluster using the variate with the highest second eigenvalue from the existing clusters, (f) iteratively apply (c), (d) and (e) until some stopping criterion is met (SAS Institute, Inc., 1985). In the present case, the procedure stopped when the largest second eigenvalue across all existing clusters fell below one.

Table 14 displays the resulting 28 thought-clusters, and Table 15 presents summary information about the clusters. While not conclusive, the cluster analysis provides further substantiation for the groupings used in analyzing the receptive fields. For the most part, thoughts cluster into easily interpretable groups, and there is substantial overlap between the groupings used in the receptive field analysis and the groupings brought out by the cluster analysis.

One of the research questions stated at the start of this chapter asks whether distributed patterns of thoughts could take the place of discrete goals in a PDP model of message production. The cluster analysis and analysis of receptive fields both provide preliminary supporting evidence for this hypothesis. Thoughts clearly seem to cluster in goal-relevant ways. Thus thoughts about cancelling cluster together (e.g., "I should stop seeing Terry for a while," "I should refuse to go with Terry tomorrow," "It would be better if Terry weren't so sure of my friendship," etc.; see cluster 22) as do thoughts about tomorrow"; see cluster 16), negative thoughts about punishing Terry (e.g., "Terry's a liar," "I want Terry to feel bad," "I want Terry to feel guilty," "I want more control in this relationship"; see cluster 8), and positive thoughts about supporting Terry (e.g., "I understand why Terry had to cancel our date," "Terry was considerate to call me," "Terry is ambitious," "Terry cares about me and my feelings," etc.; see cluster 2), to name just a few.

Some of these patterns are definitely distributed goal representations. It should be noted though, that the thought checklist contained vaguely synonymous thoughts (e.g., "I should stop seeing Terry for awhile" and "I should refuse to go with Terry tomorrow"). Theoretically, it is less interesting that these thoughts cluster than it is that conceptually related (but not synonymous) thoughts cluster. The latter sort of grouping is closer to what we mean by distributed representations.

One criticism of models that employ discrete symbolic goals (i.e., models where each goal is represented by a single atomic
primitive) is that they do not allow for continuous gradations of meaning that characterize human responses (Smolensky, 1988). This is due in part to the discrete goal's having no internal structure. Thus a discrete goal has its effects in an all-or-nothing fashion. In contrast, when goals are modelled as distributed patterns of thoughts, the goal to (for example) punish Terry can be represented in variety of slightly different ways depending on which thoughts in a goal cluster are activated and depending on how those thoughts are contextualized by other simultaneously active thoughts.

Potential for the Study of Individual Differences

In addition to providing a new model for thinking about message production, the methods described here have excellent potential for advancing the study of individual

Table 14

Thought Clusters with Conceptual Labels

Variable	R-squared with Own	R-squared with Next Closest	1-R^2 Ratio
Cluster 1 "Express anger"			
T. irritates me	.3859	.0885	.6737
I'm tired of being patient with T.	.3610	.1574	.7584
I'm upset with T.	.3396	.1232	.7532
I'm tired of T.'s behavior	.4238	.1186	.6537
T. isn't behaving like a good friend	.3941	.1777	.7368
T. repeatedly cancels appointments with me	.2631	.0969	.8159
T. thinks people should work around his/her schedule	.3290	.1822	.8204
T. needs to know that I am a little angry	.3732	.1209	.7130
Cluster 2 "Be supportive"			
I understand why T. had to cancel our date	.5266	.2033	.5942
T. was considerate to call me	.3055	.1302	.7984
T. is ambitious	.2609	.0848	.8076
T. cares about me and my feelings	.3915	.1760	.7385
The friendship is important to T.	.4994	.1277	.5739
I should be supportive to T.	.4507	.1362	.6359
T. should know that I understand about the cancellation	.5370	.3131	.6741
Cluster 3 "Make T. take responsibility for friendship"			
T. doesn't seem to have time for our friendship	.4526	.1111	.6158
T. doesn't place a high priority on our friendship	.5114	.1193	.5549
T. takes me for granted	.3445	.1774	.7968
T. should put more effort into our friendship	.5277	.1540	.5583
T. ought to take responsibility for our friendship	.3495	.1018	.7242
I should be firm in dealing with T.	.2490	.1164	.8500
Cluster 4 "T. doesn't like me anymore"			
T.'s guilty	.1142	.0454	.9279
T.'s using activities as an excuse not to see me	.4561	.0925	.5994

T. has a problem with me that he/she is afraid to bring up	.4117	.1956	.7313
T. feels our friendship is stale	.2730	.0651	.7776
T. wants to break off our friendship	.5018	.0520	.5255
T.'s avoiding me	.6442	.0575	.3775
Cluster 5 "Talk to T. about friendship"			
The friendship is important to me	.4856	.0926	.5669
Talking with T. about the friendship is important	.3776	.1473	.7299
I value T.'s friendship	.5167	.1600	.5754
T. is a close friend	.3747	.0954	.6913
T. should call me back when he/she has time for a long talk	.2970	.0528	.7422
Cluster 6 "Hide feelings"			
If I criticize T. it will hurt our friendship	.5460	.0510	.4784
If I refused to see T., T. would miss me	.0431	.0420	.9988
If T. knows I feel hurt, it will make me look bad	.6629	.0571	.3575
I don't want T. to know that I feel hurt	.5282	.0430	.4930
Cluster 7 "Make T. keep appointments"			
It's okay to cancel once in a while, but not repeatedly	.2389	.0588	.8086
T. should keep appointments	.4608	.0415	.5625
T. ought to see that cancelling appointments is bad behavior	.5655	.1075	.4868
T. needs to understand the consequences of cancelling appointments	.3700	.0580	.6688
Cluster 8 "Punish T."			
T.'s a liar	.4011	.1812	.7314
I want T. to feel bad	.7111	.1782	.3515
I want T. to feel guilty	.6234	.1487	.4424
I want more control in this relationship	.2594	.1093	.8315
Cluster 9 "T. is sincere"			
T. wants to do the movie tonight but can't because of a movie	.4633	.0828	.5851
T.'s telling the truth	.5825	.2519	.5582
T. is a leader	.1632	.0370	.8689
T. thinks it is important to keep commitments	.4582	.2050	.6815
T. feels bad about cancelling the date	.3792	.2084	.7842
T. wants me to understand what is happening in this situation	.3041	.1143	.7857
Cluster 10 "Insist on a compromise"			
T. is a happy person	.3252	.0388	.7020

If T. knew there was a problem between us, he/she would be worried	.4216	.0337	.5986
T. would respect me for being firm	.1856	.0377	.8463
T. and I could do something else and skip the movie	.5374	.0539	.4889
Cluster 11 "Make T. apologize for being a jerk"			
T.'s irresponsible	.3307	.1172	.7581
T. is selfish	.5053	.1659	.5931
T. is insecure	.1830	.0376	.8489
T. is self-centered	.5380	.2032	.5799
T. doesn't respect me	.2713	.0961	.8062
T. should apologize for being a bad friend	.3608	.1869	.7861
Cluster 12 "Insist on T. reducing other commitments"			
T. is overcommitted	.4476	.0719	.5952
T. knows too many people	.0672	.0155	.9475
T. tries to be too involved	.5808	.0754	.4533
T. expects me to be understanding	.2221	.0640	.8311
T. should make fewer commitments to others	.4484	.0456	.5780
We shouldn't make definite plans again until T. has more free time	.2384	.0678	.8169
Cluster 13 "Make T. explain his/her actions"			
If T. has a problem with me, he/she should say so	.4949	.2008	.6320
Our friendship needs to be put back on track	.4298	.0621	.6080
We need to get things out in the open	.6034	.1330	.4574
Cluster 14 "T. is gung ho"			
T. acts responsibly about attending meetings	.5711	.2177	.5483
T. is active	.2958	.1122	.7932
T. is highly motivated	.5652	.1047	.4856
T. is an overachiever	.3559	.0235	.6596
Cluster 15 "I'm disappointed"			
I was looking forward to seeing T. tonight	.4764	.0965	.5795
I feel disappointed	.4237	.0805	.6268
T. is a good friend	.4126	.0517	.6194
T. is busy	.3321	.0843	.7294
Cluster 16 "Accept request to reschedule"			
T. would like to see me tomorrow night	.5011	.1827	.6105
T. should promise not to cancel again	.1173	.0123	.8937
T. will try to keep the appointment tomorrow night	.6466	.1299	.4062

I can make it tomorrow	.5870	.1402	.4803
Cluster 17 "Make T. keep the date"			
T. should be put on the defensive	.4764	.0779	.5679
T. should cancel other commitments rather than dates with me	.3794	.0503	.6535
T. should skip the meeting tonight	.2954	.0588	.7486
I want to do to T. what he/she did to me	.4218	.1666	.6938
Cluster 18 "Go to the movie with someone else"			
This isn't the only friendship I have	.5144	.0674	.5208
I could go to the movie with someone else	.5551	.0566	.4716
I really want to see this movie	.4459	.0336	.5733
Cluster 19 "Let T. make dates"			
T. would like it if the situation were cleared up	.2749	.0335	.7503
T. thinks I want to do things with him/her	.1815	.0581	.8690
T. should be the one to make plans in the future	.4656	.0185	.5444
We should get together at a time of T.'s choosing	.6136	.0151	.3923
Cluster 20 "Tell T. how he/she has hurt me"			
T.'s actions hurt my feelings	.3735	.0937	.6913
I'm sick of being let down	.3016	.0898	.7673
T. doesn't realize that cancelling appointments can hurt others	.3303	.1248	.7651
I should let T. know how I feel about the cancellation	.3459	.1458	.7658
T. should realize what he/she has been doing to me	.3678	.1200	.7185
T. needs to know that I feel hurt	.4693	.1452	.6208
Cluster 21 "T. is insensitive"			
T. is inconsiderate	.5710	.2357	.5613
T. is insensitive to the feelings of others	.6508	.1968	.4348
T. is rude	.4237	.1848	.7070
T. doesn't consider the needs/feelings of others to be important	.5724	.1966	.5323
Cluster 22 "Refuse to reschedule"			
T. should stop acting interested in seeing me and then cancelling	.3620	.0903	.7013
I should stop seeing T. for a while	.4511	.1268	.6285
I should refuse to go with T. tomorrow	.4299	.1485	.6696
If I agree to go tomorrow, T. will probably cancel that too	.2754	.0842	.7912

I should not make dates with T. in the future	.4063	.1077	.6654
It would be better if T. weren't so sure of my friendship	.2910	.1121	.7985
Cluster 23 "Make T. explain his/her behavior"			
I don't know what T.'s motives are	.6114	.0119	.3933
T. ought to explain the repeated cancellations	.6114	.0668	.4164
Cluster 24 "Forgive T."			
I should give T. another chance	.4052	.1439	.6948
I want to spend time with T.	.6028	.1902	.4905
I want T. to go out with me	.6350	.1121	.4111
Cluster 25 "Give T. benefit of the doubt"			
T. has personal problems and obligations I don't know about	.6257	.0440	.3915
I should be understanding about T.'s problems	.6257	.1829	.4581
Cluster 26 "Make T. get organized"			
T. is disorganized	.7390	.0584	.2772
T. should organize time better	.7390	.1363	.3022
Cluster 27 "Make T. make time"			
I'm busy too, but I don't let that interfere with seeing T.	.4645	.0340	.5544
Whenever we have plans something else always comes up	.4267	.1152	.6480
T. ought to have time for me	.5123	.1450	.5704
Cluster 28			
T. should be grateful to me for being understanding	1.0000	.0461	.0000

Note. T. is an abbreviation for Terry.

differences in reasoning about messages. For example, in the analysis of receptive fields it was emphasized that substantively different (if not completely contradictory) thoughts can strongly influence the same message-frame 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, messages should only be strongly excited by a conceptually coherent set of thoughts. That this is not the case implies that message-frames are not unifunctional. Multiple (possibly contradictory) coalitions of thoughts may excite the same message-frame because that message-frame is

perceived by different sorts of people to have different sorts of functions.

This circumstance is precisely what one would expect on a design logic account of message production (Lambert & O'Keefe, 1990; O'Keefe, 1988). Lambert and O'Keefe

Table 15

Cluster Summary for 28 Thought Clusters

Cluster	Members	Cluster Variation	Variation Explained	Proportion Explained	Second Eigenvalue
1	8	8.00000	2.86964	.3587	.9376
2	7	7.00000	2.97155	.4245	.9566
3	6	6.00000	2.43466	.4058	.9192
4	6	6.00000	2.40101	.4002	.9453
5	5	5.00000	2.05164	.4103	.8559
6	4	4.00000	1.78032	.4451	.9910
7	4	4.00000	1.63527	.4088	.9331
8	4	4.00000	1.99498	.4987	.8565
9	6	6.00000	2.35048	.3917	.9897
10	4	4.00000	1.46984	.3675	.9581
11	6	6.00000	2.18904	.3648	.9001
12	6	6.00000	2.00457	.3341	.9746
13	3	3.00000	1.52817	.5094	.8191
14	4	4.00000	1.78810	.4470	.8623
15	4	4.00000	1.64484	.4112	.8705
16	4	4.00000	1.85198	.4630	.9469
17	4	4.00000	1.57293	.3932	.9450
18	3	3.00000	1.51539	.5051	.7900
19	4	4.00000	1.53565	.3839	.9972
20	6	6.00000	2.18840	.3647	.9035
21	4	4.00000	2.21786	.5545	.7242
22	6	6.00000	2.21575	.3693	.9157
23	2	2.00000	1.22285	.6114	.7771
24	3	3.00000	1.64296	.5477	.7918
25	2	2.00000	1.25141	.6257	.7486
26	2	2.00000	1.47790	.7390	.5221
27	3	3.00000	1.40347	.4678	.8308
28	1	1.00000	1.00000	1.0000	-
Total	121	121.00000	52.21067	.4315	-

(1990) argue that functions can be attributed to messages only against the background of a message design logic. To resurrect the language-game metaphor, one would say that

messages (or message-frames) are only moves: The rules of the game are what give moves their significance. Likewise, a message design logic is what defines the significance of a given content element.

What we see in the receptive fields is the combined influence of all three message design logics. Frame 54 for example ("Sorry, I can't go tomorrow"), is excited by characteristically expressive, negative thoughts about Terry's past behavior (e.g., "Terry is guilty," "Terry is rude"), by characteristically conventional thoughts about Terry's sincerity (e.g., "Terry's using activities as an excuse not to se me"), and by a characteristically rhetorical desire to reshape role relations and make a particular impression (e.g., "Terry should take responsibility for our friendship," "Terry would respect me for being firm"). By training the current model on all of the available examples, the thought-to-message mappings of expressive, conventional, and rhetorical message producers have effectively been superimposed on one another. This is why we see apparent inconsistencies and contradictions in the patterns of inhibitory and excitatory thoughts projecting to each message-frame.

The real potential of this method will best be realized by gathering sufficient examples to train three networks, where the examples are partitioned by design logic category (expressive, conventional, rhetorical). Then we would expect patterns of excitation and inhibition to reflect more unifunctionality in message-frames. That is, within a design logic, we would expect each message-frame to have a (relatively) stable significance. In addition, such a procedure would allow for the empirical testing of O'Keefe's claims about design logic by comparing the receptive fields of each message-frame across design logic categories. In this manner, the logics that connect thoughts to messages could be constructed inductively rather than being assumed (as is the case in a plan-based framework).

Summary

This chapter presented a prototype PDP model of situated message production, where situated message production is conceived as the mapping of patterns of thoughts onto patterns of message contents. The model was designed to avoid the two central problems facing plan-based 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. It was not necessary to instantiate abstract act types in the model, since thoughts map directly onto concrete message elements. Neither was it necessary to reify a single functional description of messages, since the model needs no explicit logic of action to guide selection of message contents.

The learning and generalization performance of the machine learning model was reported and analyzed. Generalization performance was found to be quite poor. A closer analysis suggested that the network needed to be trained on a larger example set in order to get better generalization performance. In spite of the poor generalization performance in cross- validation tests, a network was trained to perfection on the entire example set. A subset of the weights learned by this network was analyzed. Analysis of the weights supported the contention that message contents are multifunctional. Implications for the continued study of individual differences with these methods were discussed.

In all, the model was found to have a high degree of face validity, with weights embodying easily interpretable patterns of excitation and inhibition between thoughts and messages. Having confirmed the basic usefulness of this modelling approach for the study of message production, more examples will be gathered in support of a more extensive modelling effort.

CHAPTER EIGHT
CONCLUSION

Some words are meant to be spoken; some are better left unsaid. The point of this thesis has been to try to understand why sometimes we can't tell the difference. Anyone who has felt the healing warmth of a comforting word or the sting of careless criticism can appreciate the importance of this task.

The normal way of proceeding in this effort has been to assume that plans link goals to verbal means for achieving those goals. This idea is intuitively pleasing for several reasons, not the least of which is that it preserves the sense that messages are rational solutions to the routine dilemmas of everyday life. Unfortunately, the complexity and contingency of everyday life are transformed into uniformity and predictability by current conceptions of plan-based message production. The transformation is an inevitable consequence of trying to make abstract rationality do the work of situated practical reason (Lave, 1988; Suchman, 1987).

As long as proven alternative methods have been lacking, however, this distortion has been regarded as a necessary evil. The primary contribution of this thesis has been to offer a detailed model of situated message production that both looses the constraints imposed by the abstract rationality of planning and preserves the practical rationality of communicative action in context.

Overview

This chapter summarizes the argument against plan-based modelling. It also summarizes and interprets the findings from the three empirical investigations that comprise chapters 5, 6, and 7. The limitations of these investigations are considered, opportunities for future work are outlined, and implications are discussed.

Summary and Interpretation of Findings

Functional Abstraction and Plan-Based Models

Most message production models are based on the notion of planning as problem-solving (Newell & Simon, 1972). On this view, messages are instruments for bringing about changes in the state of the world (including changes in the intentional states of hearers, see Cohen & Perrault, 1979). A message producer begins with a representation of the desired state of the world. The desired state of the world is compared to the current state, and differences between the two states are identified. For each type of difference identified, an abstract communicative act is selected. This selection is made possible by an elaborate indexing system wherein each act is indexed by the type of difference it reduces. Abstract act types are then instantiated by concrete linguistic expressions in a process that relies on indexing decontextualized linguistic forms.

This modelling strategy depends crucially on the assumption that concrete linguistic expressions have uniform functions (i.e., they always reduce the same types of differences) by virtue of their membership in abstract categories of social action. It was the purpose of chapter 3 to demonstrate that this assumption is untenable. It is untenable because (a) there is no 'natural' linguistic unit to which category membership can be assigned (Levinson, 1979a); (b) the same expression can belong to multiple categories simultaneously (Levinson, 1979a); (c) there is no determinate procedure for assigning linguistic expressions to act type categories (Levinson, 1979a); (d) a definitive set of act types has not been and in principle

can not be constructed (D. O'Keefe, 1987, 1991), and (e) the category membership of linguistic expressions (and hence their functional significance) differs across individuals and within the same individual over the course of communicative development (O'Keefe, 1988). It was concluded from this analysis that an alternative model must avoid functional abstraction at all costs.

Message Production Without Types[48]

Chapter 4 sketched the outline of a model of situated message production that related patterns of thoughts to patterns of expressible message contents without reference to functionally indexed abstractions. The model proposes that the content selection component of Herrmann's (1983) basic model of speech production be implemented as a PDP pattern association network (Rumelhart & McClelland, 1986). In the model, both goals and messages are represented as distributed patterns of primitive elements: goals as distributed patterns of thoughts, messages as distributed patterns of expressible contents. By distributing the representation of goals over many thoughts, the system gains the capacity to represent continuous shades of goal-relevant meaning. A goal corresponds to a pattern of activation across many thoughts, and slightly different patterns represent slightly different goals. In contrast, plan-based models typically represent goals as symbolic atoms that are present or absent in all-or-nothing fashion. Furthermore, since messages are represented as patterns of concrete contents rather than as hierarchically organized abstractions, there is no instantiation problem and no need for functional indexing.

In order to implement this model of message production, it was necessary to develop new methods for eliciting thoughts and messages. More precisely, it was necessary to develop methods for representing patterns of thoughts and patterns of message elements as numeric vectors so that they could be input to a PDP learning algorithm. The next two sections describe the methods developed for eliciting and representing thoughts and messages respectively.

A Checklist Method for Thought Elicitation

The literature on verbal reports provides several techniques for eliciting people's thoughts. Interviews and open-ended questionnaires are useful, but the recording, transcription and coding of think-aloud and talk-aloud protocols is the most valid, reliable and detailed method for thought elicitation currently available (Ericsson & Simon, 1984). Considering the number of protocols that would need to be collected for initial and future modelling efforts (in the thousands), protocol techniques were deemed too labor intensive to be practical.

Instead, a pilot study was conducted to sample people's thoughts as they read and responded to the Broken Date hypothetical communication scenario (O'Keefe & Lambert, 1989). In the pilot study participants responded to open-ended probes about their thoughts, goals, expectations, fears, etc. These responses were boiled down into a checklist of 121 commonly reported thoughts. To represent the checklist responses numerically, a thought was scored as 1 if it were checked and as 0 otherwise.

The comprehensiveness and representativeness of the thought checklist was then checked against think-aloud protocols produced in response to the same situation. In all, 56% of the thoughts reported in the think-aloud protocols were coded as already present on the thought checklist. A small number of possible additions to the checklist were suggested, but based on the degree of overlap between methods of elicitation, the checklist was concluded to be an adequate measure of people's thoughts about the Broken Date scenario. In fact, based on subsequent analysis, there appears to be substantial redundancy among the

[48] The title of this section is a slight play on Sibun's (1991) paper entitled "Generating text without trees".

elements of the checklist. Perhaps it could be improved by removing redundant elements and extremely low-frequency elements. This is a task for future research.

Semi-Automated Message Coding

Having developed a method for rendering patterns of thoughts into a numeric vector representation, the next step was to develop a method for coding messages into an equivalent format. Because the open-ended nature of the message elicitation method is crucial to its validity, a message checklist analogous to the thought checklist was not considered.

Messages are typically thought to have hierarchical structure. Vectors, on the other hand, are inherently 'flat'. The problem was to represent these structured objects with binary features. One way to accomplish this feat is to assign a binary feature to every possible structure in a domain and allow each feature to represent the presence or absence of a single structure. This strategy fails when the number of structures or structural variations in a domain makes the required feature vectors grow to unreasonable lengths. The purpose of chapter 6 was to (a) analyze the extent of structural variation in messages (actually message elements), (b) to determine whether the flat featural representation would be practical, and (if so) (c) to develop a coding system that would translate messages into feature vectors.

As it turned out, when messages were segmented into independent clauses, it was possible to identify a (relatively) small number (72) of content categories into which all of the message elements could be coded. Each content category was considered to be a message frame: an idea-complex expressed by alternative syntactic structures in a common vocabulary.

The most significant finding of chapter 6 was the discovery of substantial structural similarity and redundancy among message elements within the same frame. This was interpreted as evidence of phrasal-lexical processing (see e.g., Becker, 1975; Dyer, 1989). Contrary to standard images of language production, wherein phrases are constructed piece by piece (Levelt, 1989), message production with a phrasal lexicon involves the selection of preformulated phrases from a existing set (in much the same way that standard models select words from a mental lexicon). Significantly, phrasal-lexical models of message production need not resort to functional indexing because abstract types are never instantiated; phrase tokens are simply selected and output as wholes.

With 72 message-frames and thousands of message elements to code, coding reliability became a concern. The demands of a 72 category coding system overwhelm human short-term memory limitations. Worse still, standard reliability measures (e.g., Cohen's kappa, [1960]) assume that members of each category occur frequently in each reliability sample. With 72 categories and several relatively infrequently occurring frames, this assumption is not met. Therefore, a decision was made to automate the coding system (as much as possible) in order to satisfy reliability concerns. Others have adopted the same strategy in similar circumstances (Bhaskar & Simon, 1977; Waterman & Newell, 1971, 1972, 1973).

A prototype automatic coding system (AutoCoder) was developed as a LISP computer program. The performance of AutoCoder on novel messages was assessed and found to be limited. The system achieved a complete and correct classification for only 16% of the novel message elements it attempted, though the rate of partial classification was much higher.

An analysis of misclassifications suggested that performance could be improved by adding more exemplars to the coding system and relaxing the system's exact-matching criterion. The result of these modifications may result in a semi-automated rather than a fully-automated coding system, but this would be sufficient. After all, the goal is to address reliability concerns, not to build an automatic coding system for its own sake. Human codings of message elements preprocessed by AutoCoder were sufficiently reliable for the present purposes.

Mapping Thoughts onto Messages

The penultimate chapter (chapter 7) described a PDP model that maps patterns of thoughts onto patterns of message frames using a standard, feed-forward, pattern association learning paradigm (Fahlman, 1988; Rumelhart & McClelland, 1986). The network's generalization performance (i.e., its predictive accuracy) was evaluated by 10-fold cross-validation. For novel cases, the network correctly predicted approximately 88% of the bits in the output feature vector. However, since most of the output features are off most of the time, 88% was an overly optimistic estimate. When performance was assessed in terms of on-hits (the number of output elements that were turned on when they were supposed to be turned on), generalization fell considerably to 22.1% correct.

Closer analysis of the relationship between frame frequency and per-frame generalization performance revealed that better performance was associated with more frequently occurring frames. Since a network can not extract useful generalizations from only a few examples, this frequency effect was not surprising. Additional analyses showed a strong negative correlation between frame frequency and off-hit generalization performance and between frame frequency and overall performance. On the surface this would seem to indicate that any improved on-hit performance would be offset by degraded off-hit performance when more examples were included in the training set. On the contrary, the findings were interpreted as evidence of a U-shaped curve relating performance to frame frequency: performance peaks with very few and very many examples. It was concluded from the analysis of generalization performance that more examples should be gathered. Nevertheless, for the purposes of establishing the face validity of the model, further analysis of network behavior proceeded.

Having evaluated the generalization performance of the via cross-validation, the network was subsequently trained to perfection on all available data. The weight matrix learned during this training was analyzed in detail. The purpose of detailed analysis of learned weights was to begin to unpack the "logic" that connects thoughts to messages. In a planning model, the function of a message (and hence the logic of message design) is assumed a priori and built into the structure of the model by means of functional indexing. In the situated, phrasal-lexical model, no assumptions are made about message function. Instead, an attempt is made to reconstruct the logic that connects thoughts to messages by examining the nature of the learned thought-message mapping.

In order to understand the relationship between message-frames and thoughts, the excitatory and inhibitory receptive fields of the five most accurately predicted message-frames were examined. The 12 thoughts with the strongest positive weights to a given unit comprised that unit's excitatory receptive field. The 12 thoughts with the strongest negative weights comprised the unit's inhibitory receptive field. By looking at which thoughts strongly excite and inhibit a given frame, it is possible to get a sense of the frame's functional significance in the Broken Date context.

In chapters 3 and 4 it was asserted that message elements are multifunctional, and this is precisely what the receptive fields indicated. For each frame analyzed, a few distinct coalitions of thoughts acted to excite and inhibit the frame in question. For the most part, each coalition of thoughts was internally coherent but conceptually distinct from the others in the receptive field. Separate coalitions were, in many cases, straightforwardly interpretable in terms of expressive, conventional, or rhetorical reasoning (O'Keefe, 1988). The multifunctionality of each frame was illustrated by the distinctive (and sometimes contradictory) kinds of thought coalitions that were found simultaneously to excite the same frame. If frames were truly unifunctional, as planning models assume, each frame should

have been excited by a single conceptually coherent set of thoughts. This was not the case. Instead, conceptually distinct sets of thoughts were simultaneously excitatory (or inhibitory). This was interpreted as evidence that each frame had multiple potential functions.

To support the claim that thoughts tended to occur together in specific patterns, the thought patterns were subjected to an oblique principal component cluster analysis. Twenty-eight clusters were formed by this procedure. Several of the clusters identified by this procedure mirrored those which were grouped together solely on conceptual grounds as part of the analysis of receptive fields. Thus, the cluster analysis was successful in providing (partial) corroboration for the thought groupings used in the receptive field analysis. Not only were the specified thoughts conceptually related, but there is evidence to indicate they occurred together as well. Taken together, this pattern of findings lends credence to the notion that, within a model of situated message production, goals can be conceived of as distributed patterns of situation-specific thoughts. Furthermore, a single message-frame may be multifunctional and simultaneously indexed by multiple (perhaps contradictory) goals.

Limitations of the Current Investigation

Nearly all of the significant limitations of this thesis are a consequence of its preliminary nature. Particular limitations will be addressed in detail below, however, the common origin of the limitations must be kept in mind.

The empirical portion of this thesis was an attempt to devise and implement a new model of message production based on principles fundamentally different from those of existing models. To the extent that the model addressed novel questions and concerns, methods needed to be devised that could bear a novel sort of evidentiary burden. Similarly, to the extent that previous research was criticized for its methodological errors (or its operationalizations of untenable theoretical assumptions), new methods and operationalizations needed to be devised that would avoid these errors.

The current thesis is an initial exploration of a set of new ideas. As such the goals were to establish the feasibility, plausibility, and/or face validity of new procedures rather than to provide definitive answers to the deepest underlying questions. The point is not to excuse the limitations of the current project but to establish a reasonable context for the interpretation of their significance.

Limitations of the Thought Checklist Task

Though chapter 5 demonstrated the basic adequacy of the thought checklist as a measure of people's thoughts, the checklist method still has its limitations. On one hand, it is not an exhaustive list of possible thoughts about the Broken Date situation. As the comparison with think-aloud protocols illustrates, there are some commonly reported thoughts that the checklist does not include. Thoughts that were relatively common, interesting from a theoretical point of view, and truly novel (in the sense that the checklist contained no synonymous thoughts), were suggested as possible additions to future versions of the thought checklist. On the other hand, the checklist is very long (121 thoughts); it contains redundant elements; and some thoughts are checked only very rarely. Therefore, the addition of new entries to the checklist should be accompanied by the removal of redundant thoughts and extremely low frequency thoughts.

Limitations of the Message Frame Coding System

The main limitation of the AutoCoder system is its poor ability to correctly and completely classify previously unseen message elements. In initial tests, the system correctly coded only 16% of the novel message elements attempted. This level of performance did not sufficiently address the coding reliability concerns that motivated the system's construction in the first place. However, analysis of coding errors indicated that the system performed better on frames for which it had stored more exemplars. If the coding rules for less common

message frames were elaborated, overall performance would improve. Other additions should also be made. As it stands, the system only recognizes exact matching as success. This overly strict criterion should be relaxed to admit partial matches and probabilistic classification.

Next, the system relies completely on expert human coders for its knowledge. Expert coders must supply all phrase and keyword definitions as well as target categories. The process of developing these coding rules is extremely time-consuming and tedious. It would be best if the system had some capacity for finding its own patterns in the messages and inducing its own coding rules from correctly coded examples.

The system's coding speed is another limitation. Search through the space of possible phrases is now done serially because there were too many options in the phrase definitions to expand them all into a tree that could be searched in parallel. In addition, there are limitations in the representation. Broadly speaking, the representational system for message-elements, keywords and phrases is inconsistent and somewhat ad hoc. For example, keyword definitions and phrase definitions are in stored in slightly different formats. Where possible, the representation should be consistent across levels of abstraction. In another case, keywords defined purely in terms of primitive words are stored in the same data structure as keywords defined in terms of other keywords. Logically, there should be two data structures: one for first-order definitions, one for second-order definitions. This change would increase coding speed because search could be directed to one of the two data structures at the appropriate time.

Limitations of the PDP Model

The PDP model mapping thoughts to messages suffers from several limitations. First of all, the PDP model exhibits poor generalization performance, correctly activating only 22% of the message-frames it was supposed to activate in response to previously unseen thought patterns. Analysis of patterns of misclassification indicates that more frequently occurring frames were more accurately predicted. Thus, the system needs to learn form a larger example set before its generalization performance will improve. Due to the poor generalization performance, only part of the learned weight matrix was analyzed in detail. A more complete analysis is necessary.

Representational decisions also limited the current modelling project. To begin with, the sequential placement of message frames within a message was ignored. If the same frame occurred more than once in a message, all but the first occurrence was ignored. This is a distortion of the raw data that needs to be overcome. Beyond this distortion, the sequence of message-elements was represented as a concurrent pattern of thoughts, and there was no mechanism offered for recovering the sequence from the pattern. One might suggest that the frames be output in decreasing order of their activations, with the most active output first and least active output last. A more sophisticated solution would employ a PDP architecture capable of producing sequential output (see, e.g., Elman, 1990).

Finally, the current model captures the thought-message mapping at a single point in time, but it offers no account of how the mapping changes over the course of development. This is an especially acute weakness considering that the constructivist (O'Keefe & Delia, 1982, 1985) and message design logic (O'Keefe, 1988) traditions on which this work is based both seek to understand the development of functional communication skills over the life-span. The following section on future work includes a description of one possible framework for modelling developmental phenomena.

Future Work

This thesis established the potential of several new methods for studying and modelling message production. A great deal of work must now be done to fulfill that potential.

Initially, the methods developed and tested here should be refined. The thought checklist should be revised, pruning redundant thoughts and adding novel ones. The AutoCoder coding system should be modified in the manner suggested above, with more coding rules added to cover less frequently occurring message frames. The possibility of building a version of AutoCoder that learns its own coding rules from correctly coded examples should be explored. Currently available syntactic pattern recognition algorithms and machine learning techniques are suitable for this task.

Beyond these methodological refinements, the top priority of future work should be to retrain the PDP model on a scale large enough to significantly improve generalization. At least an order of magnitude more examples (N = 3000) should be gathered. Once a suitable level of generalization performance is achieved (say greater than 75% on-hits for novel cases), the entire learned weight matrix could be examined in detail. Any of a variety of multivariate methods (see, e.g., Hanson & Burr, 1990) could then be used to analyze the networks response to expressive, conventional, and rhetorical messages respectively. Such an analysis would provide evidence bearing on O'Keefe's (1988) claims about individual differences in reasoning about communication.

The next major step would be to model developmental processes. A lengthy discussion is necessary in order to explain how to take the next step. The image of message production and the development of competence described below is an extension of ideas first articulated by Rumelhart and Jordan (Jordan & Rosenbaum, 1989; Jordan & Rumelhart, 1990; Rumelhart, Smolensky, McClelland, & Hinton, 1986).

Thus far an attempt has been made to map from current and desired state (represented as a single pattern of thoughts) to a message. Theoretically, the message is chosen because its utterance is expected to transform the current state into the goal state. In symbolic terms:

(21) $S(I_c, I_g) = U_{c,g}$

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 world model E():

(22) $E_c(I_c, U_{c,g}) = I_g$

In words, E() takes a message and the current state as input and outputs a description of its effect. With (21) and (22) in mind, a more complete process-story can be told about message production and about development.

I will refer to (21) as the message generator and (22) as the world model. A complete model of message production includes both (see Figure 8). In theory they are connected, with the output of the message generator being the input to world model:

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

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

(24) $S(I_g) \rightarrow U_g \rightarrow E(U_g) = I_g$

Within this model of message production, developmental phenomena can be neatly conceptualized.

Imagine that the message generator and the world model are both implemented as PDP networks. Early in life, the weights in the message generator connecting goals (patterns of thoughts) to messages are random, as are the weights connecting messages to effects in the world model. During this early period, message production proceeds as follows: While engaged in some activity, a child develops a communicative goal. This goal maps to a message by way of the message generator. The message is evaluated by the world model to determine whether its utterance brings about the desired effect; if so the message is uttered. If not, the model of effects propagates an error message back through the message generator, changing the message generator's weights. Then another candidate message is generated, evaluated, and the process repeats until the child is satisfied with the chosen message. Notice though, that since the message generator and the model of effects are initially randomly wired, message selection is uninformed, and evaluation is inaccurate.

As a child gains experience in the social world, however, he/she becomes progressively more adept at predicting the effects of messages in context. This learning corresponds to modification of the weights connecting messages to their effects. The improved ability to anticipate effects translates into better feedback to the message generator, and thus to better initial message selections.

As a person matures, the world model begins to change less, and explicit evaluation of generated messages becomes less and less necessary because, through extended feedback, the world model is essentially reflected in the connections between distributed goals and messages. The model of effects is <u>not</u> internalized by the message generator. Rather, the message generator organizes goal-message relations with respect to a given model of effects. Explicit evaluation is only necessary in novel or extraordinarily difficult situations (where effects are hard to predict or accurate prediction is crucial).

With this comprehensive model in mind, it is possible to get a better perspective on what has been done and on what is left to do. The PDP model of chapter 7 is essentially a message generator, trained without the benefit of a stable, specific model of effects.[49] Instead, the thought-message relations embodied by the learned weights analyzed in chapter 7 reflect the influence of many different models of effects, each from an individual at a different stage of communicative development.

The task ahead is to figure out how effects (and hence goal-states) are represented. Which of the thoughts in the current thought checklist correspond to the representation of the current state and which to the goal state? Perhaps only thoughts that predicate future states of self, Terry, the friendship, etc. are representative of the goal state. If this full-blown model of the development of communicative competence is ever to be implemented, future work must address this issue.

To implement this model, examples need to be gathered to train the world model. The input half of these examples will be messages (expressed as 72-bit message-frame vectors), and the output half will be effects (goal states likely represented by some subset of the current thought checklist). It is not immediately obvious how to collect these examples. Perhaps subjects could be given a description of the Broken Date scenario and a message; then they would be asked to check from the goal-state checklist those statements that best described the situation after the message had been spoken. The development of a system for eliciting these examples in the required format is an important task for future research.

If it were possible to gather such examples, one very interesting modelling project

[49] This corresponds to Jordan and Rumelhart's (1990) "direct inverse model."

would be possible. The examples could be partitioned by design logic category, and three separate models of message effects could be trained (corresponding to expressive, conventional, and rhetorical design logics respectively). Then these models of message effects could be used as the source of error feedback to train three separate message generators (see Jordan & Rumelhart, 1990, for details of this training process). This procedure would, in effect, "grow" one expressive, one conventional, and one rhetorical message generator.

Implications and General Discussion

The most significant theme developed in this thesis concerns the misuse of analysts' abstractions, particularly the reification of functional abstractions in models of language production. By briefly reviewing the main argument and main contribution of the thesis, I will once more try to expand upon this theme.

Clearly one message to take away from this exercise is that plan-based models of message production need to be carefully reevaluated. First, since they operate on functional abstractions such as act types (Appelt, 1986) or pragmatic or rhetorical features (Hovy, 1990), and since these abstractions must eventually be instantiated, plan-based models must employ functional indexing schemes. However, the very idea of such indexing is unsound. Indexing fails (a) because decontextualized linguistic forms have no functional significance, and (b) because contextualized means-ends relations are not uniform across or within individuals. These are not problems with planning models per se but rather with planning models applied to language production. Specifically, means-ends analysis requires that operators be indexed by the differences they reduce, but this strategy is workable only when operators have stable, context-independent effects. When this requirement is not met, and it is not met in the domain of language, models that rely on explicit indexing are certain to prove faulty.

Second, the reification of analysts' abstractions in computational or theoretical models is equally harmful when concrete situations (rather than messages) are recast as types of situations or situations embodying types of goals. This transformation runs counter to the arguments that message production is a situated activity the performance of which is tied inextricably to the details of a situation (see e.g., Lambert & O'Keefe, 1989; Lave, 1988; Suchman, 1987). Because the act of abstraction strips away the details of a particular situation, resulting situation- or goal-types are not representationally rich enough to allow for the context-sensitive indexing of messages. Third, in spite of overwhelming evidence of diversity in linguistic means-ends relations, plan-based models build in a fixed set of situation-type/message-type associations.

There is a lesson here for models of language production: do not reify analytical abstractions by giving them a causal role in cognitive processing. The most serious problems of plan-based models are all related to the reification of analysts' abstractions. The introduction of types into a computational model either depletes information from the representation of situations (in the case of situation-types), introduces an instantiation problem for which no adequate solution exists (in the case of message-types), or prevents the appearance of message diversity (in the case of means-ends types). The only sensible alternative is to model message production without reifying types. This is the line of reasoning that led to the development of the current content-based model.

Content-Based Models of Message Production

This analysis of the role of abstraction and the problems of indexing indicates a need to move away from models that first generate descriptions of messages and then instantiate them. Rather than generate-then-instantiate, models of message production should be based on a generate-then-evaluate flow of control as described above (see Jordan & Rosenbaum,

1989; Jordan & Rumelhart, 1990). Such models are content-based rather than act-based because they operate solely at the level of instantiated content; they model goals as distributed patterns of situation-specific thoughts, and they generate message contents rather than descriptions of message contents.

Still, the intuition that production is goal driven is preserved in the evaluation stage (see Figure 8), where the change of state caused by the expression of selected content is predicted and compared to the desired change of state. In addition, all of the processing in content-based, generate-then-evaluate models is contextualized. Current and goal states are patterns of situation specific thoughts, and message effects are modelled as a function of message contents and the thoughts that comprise the current state. Thus messages are only associated with effects in a specific context.

Content-based models are attractive at this stage primarily because of the mistakes they avoid. If they are to prove useful in the long run, they must offer positive contributions of their own. At present, they seem to suggest new and interesting ways of explaining individual differences and the development of communication competence without blatantly contradicting established evidence, and while providing a way around the instantiation problem. Even granting this limited success, only continued development and evaluation of content-based models such as the one proposed here will decide whether they represent a truly favorable alternative to existing plan-based models.

REFERENCES

Agre, P. E. (1990). Review of Plans and situated actions: The problem of human-machine Communication. Artificial Intelligence, 43, 369-384.

Agre, P. E. (in press). The dynamic structure of everyday life. Cambridge: Cambridge University Press.

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.

Allen, J. F. (1983). Recognizing intentions in natural language utterances. In M. Brady & R. C. Berwick (Eds.), Computational models of discourse (pp. 107-166). Cambridge, MA: MIT Press.

Allen, J. F., & Perrault, C. R. (1986). Analyzing intention in utterances. In B. J. Groscz, K. Sparck-Jones, & B. L. Webber (Eds.), Readings in natural language processing (pp. 441-458). Los Altos, CA: Morgan Kaufmann.

Anderson, J. R. (1976). Language, memory, and thought. Hillsdale, NJ: Lawrence Erlbaum Associates.

Anderson, J. R. (1983). The architecture of cognition. Cambridge, MA: Harvard University Press.

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

Appelt, D. E. (1986). Planning English referring expressions. In B. J. Grosz, K. Sparck-Jones, & B. L. Webber (Eds.), Readings in natural language processing (pp. 501-517). Los Altos, CA: Morgan Kaufmann.

Applegate, J. L., & Delia, J. G. (1980). Person-centered speech, psychological development, and the contexts of language usage. In R. St. Clair & H. Giles (Eds.), The social and psychological contexts of language (pp. 245-282). Hillsdale, NJ: Lawrence Erlbaum Associates.

Austin, J. L. (1962). How to do things with words. New York: Oxford University Press.

Bach, K., & Harnish, R. M. (1979). Linguistic communication and speech acts. Cambridge, MA: MIT Press.

Barr, A., & Feigenbaum, E. A. (1981). The handbook of artificial intelligence: Vols. 1-2. Los Altos, CA: William Kaufmann, Inc.

Becker, J. (1975). The phrasal lexicon (Artificial Intelligence Report No. 28). Cambridge, MA: Bolt Beranek & Newman Inc.

Berger, C. (1987). Planning and scheming: Strategies for initiating relationships. In R. Burnett, P. McGhee, & D. Clarke (Eds.), Accounting for relationships: Social representations of interpersonal links (pp. 158-174). London: Methuen.

Berger, C. (1988). Planning, affect and social affect generation. In R. L. Donohew, H. Sypher, & E. T. Higgins (Eds.), Communication, social cognition, and affect (pp. 93-

116). Hillsdale, NJ: Lawrence Erlbaum.
Berger, C., & Bell, R. (1988). Plans and the initiation of social relationships. Human Communication Research, 15, 217-235.
Bhaskar, R., & Simon, H. A. (1977). Problem solving in semantically rich domains: An example from engineering thermodynamics. Cognitive Science, 1, 193-215.
Blumer, H. (1969). Symbolic interactionism: Perspective and method. Englewood Cliffs, NJ: Prentice-Hall.
Bock, J. K. (1982). Towards a cognitive psychology of syntax: Information processing contributions to sentence formulation. Psychological Review, 49, 1-47.
Boden, M. A. (1988). Computer models of mind. Cambridge: Cambridge University Press.
Brown, P., & Levinson, S. C. (1987). Politeness: Some universals in language usage. Cambridge: Cambridge University Press.
Bruce, B. C. (1980). Plans and social actions. In R. J. Spiro, B. C. Bruce, & W. F. Brewer (Eds.), Theoretical issues in reading comprehension (pp. 367-384). Hillsdale, NJ: Lawrence Erlbaum Associates.
Bruce, B. C. (1986). Generation as social action. In B. J. Groscz, K. Sparck-Jones, & B. L. Webber (Eds.), Readings in natural language processing (pp. 419-422). Los Altos, CA: Morgan Kaufmann.
Burleson, B. R. (1984). Comforting communication. In H. E. Sypher & J. L. Applegate (Eds.), Communication by children and adults: Social, cognitive, and strategic processes (pp. 63-105). Newbury Park, CA: Sage.
Burleson, B. R., & Samter, W. (1985). Consistencies in theoretical and naive evaluations of comforting messages. Communication Monographs, 52, 103-125.
Butterworth, B. (Ed.). (1980). Language production: Vol. 1, Speech and talk. London: Academic Press.
Chapman, D. (1987). Planning for conjunctive goals. Artificial Intelligence, 32, 333-377.
Clark, A. (1989). Microcognition: Philosophy, cognitive science and parallel distributed processing. Cambridge, MA: MIT Press.
Clark, H. H., & Clark, E. V. (1977). Psychology and language. New York: Harcourt Brace Jovanovich.
Chafe, W. (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 Press.
Cohen, J. (1960). A coefficient of agreement for nominal scales. Education and Psychological Measurement, 20, 37-46.
Cohen, J., & Cohen, P. (1983). Applied multiple regression/ correlation analysis for the behavioral sciences. Hillsdale, NJ: Lawrence Erlbaum Associates.
Cohen, P. R., & Feigenbaum, E. A. (1981). The handbook of artificial intelligence: Vol. 3. Los Altos, CA: William Kaufmann, Inc.
Cohen, P. R., Morgan, J., & Pollack, M. E. (Eds.). (1990). Intentions in communication. Cambridge, MA: MIT Press.
Cohen, P. R., & Levesque, H. J. (1990). Rational interaction as the basis for communication. In P. R. Cohen, J. Morgan, & Pollack, M. E. (Eds.), Intentions in communication (pp. 221-255). Cambridge, MA: MIT Press.
Cohen, P. R., & Perrault, C. R. (1979). Elements of a plan-based theory of speech acts. Cognitive Science, 3, 177-212.
Cooper, C. & Odell, B. (Eds.). (1978). Research on composing: Points of departure. Urbana, IL: National Council of Teachers of English.
Craig, R. T. (1989). Communication as a practical discipline. In B. Dervin, L. Grossberg, B.

O'Keefe, & E. Wartella (Eds.), Rethinking communication, Vol. 1: Paradigm issues (pp. 97-122). Newbury Park, CA: Sage.
Deese, J. E. (1984). Thought into speech: The psychology of a language. Englewood Cliffs, NJ: Prentice-Hall.
Delia, J. G., Kline, S., & Burleson, B. (1979). The development of persusaive communication strategies in kindergarteners through twelfth-graders. Communication Monographs, 46, 241-256.
Dell, G. S. (1986). A spreading activation theory of retrieval in sentence production. Psychological Review, 93, 283-321.
Dell, G. S. (1988). The retrieval of phonological forms in production: Tests of predictions from a connectionist model. Journal of Memory and Language, 27, 124-142.
Dell, G. S., & Reich, P. A. (1981). Stages in sentence production: An analysis of speech error data. Journal of Verbal Learning and Verbal Behavior, 20, 611-629.
Dell, G. S., & Juliano, C. (1990). Connectionist approaches to the production of words. Unpublished manuscript, Department of Psychology, University of Illinois, Urbana, IL.
Dennett, D. C. (1981). Brainstorms. Cambridge, MA: MIT Press.
Dennett, D. C. (1987). The intentional stance. Cambridge, MA: MIT Press.
van Dijk, T. A., & Kintsch, W. (1983). Strategies of discourse comprehension. New York: Academic Press.
Dillard, J. P. (Ed.). (1990). Seeking compliance: The production of interpersonal influence messages. Scottsdale, AZ: Gorsuch Scarisbrick.
Dyer, M. G. (1989). Comprehension and acquisition of figurative expressions with phrasal/lexical memory. Metaphor and Symbolic Activity, 4, 173-201.
Elman, J. L. (1990). Finding structure in time. Cognitive Science, 14, 179-211.
Ericsson, K. A., & Simon, H. A. (1984). Protocol analysis: Analyzing verbal reports as data. Cambridge, MA, MIT Press.
Ernst, G. (1987). Means-ends analysis. In S. C. Shapiro, D. Eckroth, & G. A. Vallasi (Eds.), Encyclopedia of artificial intelligence (pp.578-584). New York: John Wiley & Sons.
Farah, M. (1988). Is visual imagery really visual? Overlooked data from neuropsychology. Psychological Review, 95, 307-317.
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.
Feldman, J. A., & Ballard, D. H. (1982). Connectionist models and their properties. Cognitive Science, 6, 205-254.
Fikes, R. E., & Nilsson, N. J. (1971). STRIPS: A new approach to the application of theorem proving to problem solving. Artificial Intelligence, 2, 189-208.
Fillmore, C. J. (1968). The case for case. In E. Bach & R. Harms (Eds.), Universals in linguistic theory (pp. 1-90). New York: Holt, Rinehart and Winston.
Flavell, J. H., Botkin, P. T., Fry, C. L., Wright, J. W., & Jarvis, P. E. (1968). The development of role-taking and commuication skills in children. New York: Wiley.
Fodor, J. A., & Pylyshyn, Z. (1988). Connectionism and cognitive architecture: A critical analysis. Cognition, 28, 3-71.
Garfield, J. L. (Ed.). (1987). Modularity in knowledge representation and natural-language processing. Cambridge, MA: MIT Press.
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. (1980). Levels of processing in sentence production. In B. Butterworth (Ed.), Language production: Vol. 1. Speech and talk (pp. 177-220). London: 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 University Press.
Garrett, M. F. (1990). Sentence processing. In D. N. Osherson & H. Lasnik (Eds.), An invitation to cognitive science: Language, Vol. 1 (pp. 133-175). Cambridge, MA: MIT Press.
Gazdar, G. (1981). Speech act assignment. In A. K. Joshi, B. L. Webber, & I. A. Sag (Eds.), Elements of discourse understanding (pp. 64-83). Cambridge: Cambridge University Press.
Glass, G. V., & Hopkins, K. D. (1984). Statistical methods in education and psychology. Englewood Cliffs, NJ: Prentice Hall.
Goffman, E. (1967). Interaction ritual: Essays on face-to-face behavior. New York: Pantheon.
Goffman, E. (1981). Forms of talk. Philadelphia, PA: University of Pennsylvania Press.
Graesser, A. C., & Bower, G. H. (Eds.). (1990). Inferences and text comprehension. New York: Academic Press.
Greene, J. O. (1984a). A cognitive approach to human communication: An action assembly theory. Communication Monographs, 51, 289-306.
Greene, J. O. (1984b). Speech preparation processes and verbal fluency. Human Communication Research, 11, 61-84.
Grice, H. P. (1975). Logic and conversation. In P. Cole & J. Morgan (Eds.), Syntax and semantics 3: Speech acts (pp. 41-58). New York: Academic Press.
Grosz, B. J., Sparck-Jones, K., & Webber, B. L. (1986). Readings in natural language processing. Los Altos, CA: Morgan Kaufmann.
Guetzkow, H. (1950). Unitizing and categorizing problems in coding qualitative data. Journal of Clinical Psychology, 6, 47-58.
Hanson, S. J., & Burr, D. J. (1990). What connectionist models learn: Learning and representation in connectionist networks. Behavioral and Brain Sciences, 13, 471-518.
Hanson, S. J., & Kegl, J. (1987). PARSNIP: A connectionist model that learns natural language grammar from exposure to natural language sentences. In Proceedings of the Ninth Annual Cognitive Science Conference, Seattle, WA.
Harmon, H. H. (1976). Modern factor analysis. Chicago, IL: University of Chicago Press.
Hayes-Roth, B., & Hayes-Roth, F. (1979). A cognitive model of planning. Cognitive Science, 3, 275-310.
Hobbs, J., & Evans, D. (1980). Conversation as planned behavior. Cognitive Science, 4, 349-377.
Heritage, J. (1984). Garfinkel and ethnomethodology. Cambridge: Polity Press.
Herrmann, T. (1983). Speech and situation: A psychological conception of situated speaking. Berlin: Springer-Verlag.
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. H. (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. Beverly Hills, CA: Sage.
Johnson-Laird, P. N. (1983). Mental models. Cambridge: Cambridge University Press.

Jordan, M, & Rosenbaum, D. (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.
Kandel, E. R., & Schwartz, J. H. (1985). Principles of neural science. New York: Elsevier/North Holland.
Keller, C., & Keller, J. (1988). Thinking and acting with iron. Unpublished manuscript, Department of Anthropology, University of Illinois, Urbana.
Kellerman, K., Broetzmann, S., Lim, T., & Kitao, K. (1989). The conversation MOP: Scenes in the stream of discourse. Discourse Processes, 12, 27-61.
Kempen, G. (Ed.). (1987). Natural language generation. Dordrecht, Holland: Martinus Nijhoff.
Kosslyn, S. M., & Pomerantz, J. R. (1977). Imagery, propositions, and the form of mental representations. Cognitive Psychology, 9, 52-76.
Lachenbruch, P. A., & Mickey, M. R. (1968). Estimation of error rates in discriminant analysis. Technometrics, 10, 1-11.
Lachter, J., & Bever, T. (1988). The relations between linguistic structure and associative theories of language learning: A consructive critique of some connectionist leatning models. Cognition, 28, 195-247.
Lambert, B. L. (1991). AutoCoder [Computer program]. Urbana, IL: University of Illinois.
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.
Langley, P., Simon, H. A., Bradshaw, G. L., & Zytgow, J. M. (1987). Scientific discovery: Computational explorations of the creative processes. Cambridge, MA: MIT Press.
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.
Levine, D. S. (1991). Introduction to neural and cognitive modeling. Hillsdale, NJ: Lawrence Erlbaum Associates.
Levinson, S. C. (1979a). Some pre-observations on the modelling of dialogue. Discourse Processes, 4, 93-116.
Levinson, S. C. (1979b). Activity types and language. Linguistics, 17, 365-399.
Levinson, S. C. (1983). Pragmatics. Cambridge: Cambridge University Press.
Levy, D. M. (1979). Communicative goals and strategies: Between discourse and syntax. In T. Givon (Ed.), Syntax and semantics 12: Discourse and syntax (pp. 183-212). New York: Academic Press.
MacWhinney, B., Leinbach, J., Taraban, R., & McDonald, J. (1989). Language learning: Cues or rules. Journal of Memory and Language, 28, 255-277.
McClelland, J. L., & Kawamoto, A. H. (1986). Mechanisms of sentence processing: Assigning roles to constituents of sentences. In J. L. McClelland & D. E. Rumelhart (Eds.), Parallel distributed processing: Explorations in the microstructure of cognition, Vol. 2: Psychological and biological models (pp. 272-325). Cambridge, MA: MIT Press.
McClelland, J. L., & Rumelhart, D. E. (1986). Parallel distributed processing: Explorations in the microstructure of cognition, Vol. 2: Psychological and biological models. Cambridge, MA: MIT Press.
McDonald, D. D. (1987). Natural-language generation. In S. Shapiro, D. Eckroth, & G. Vallasi (Eds.), Encyclopedia of artificial intelligence, Vol. 1 (pp. 642-655). New York: John Wiley & Sons.

McDonald, D. D., & Bolc, L. (Eds.). (1988). Natural language generation systems. New York: Springer-Verlag.
McKeown, K. R., & Swartout, W. R. (1988). Language generation and explanation. In M. Zock & G. Sabah (Eds.), Advances in natural language generation: An interdisciplinary perspective (1-51). London: Pinter.
Miller, D. (1990). Towards a believable theory of planning [Multiple review of Practical planning, Blueprints for thinking, and Case-based planning]. Journal of Mathematical Psychology, 34, 489-498.
Miller, G., Galanter, E., & Pribram, K. (1960). Plans and the structure of behavior. New York: Holt, Rhinehart and Winston.
Moore, R. C. (1980). Reasoning about knowledge and action. SRI International Center Technical Note No. 191. Menlo Park, CA.
Morrow, D. G., Bower, G. H., & Greenspan, S. L. (1989). Updating situation models during narrative comprehension. Journal of Memory and Language, 28, 292-312.
Morrow, D. G., Greenspan, S. L., & Bower, G. H. (1987). Accessibility and situation models in narrative comprehension. Journal of Memory and Language, 26, 165-187.
Newell, A. (1973). Production systems: Models of control structure. In W. Chase (Ed.), Visual information processing (pp. 463-526). New York: Academic Press.
Newell, A. (1980). Physical symbol systems. Cognitive Science, 4, 135-183.
Newell, A., Shaw, J. C., & Simon, H. A. (1963). Empirical explorations with the logic theory machine. In E. A. Feigenbaum & J. Feldman (Eds.), Computers and thought (pp. 109-133).
Newell, A., & Simon, H. A. (1972). Human problem solving. Englewood Cliffs, NJ: Prentice-Hall.
Newell, S., & Stutman, R. (1982, November). A qualitative approach to social confrontation: Identification of constraints and facilitators. Paper presented at the annual meeting of the Speech Communication Association, Louisville, KY.
O'Keefe, B. J. (1988). The logic of message design: Individual differences in reasoning about communication. Communication Monographs, 55, 80-103.
O'Keefe, B. J. (1990a). The logic of regulative communication: Understanding the rationality of message designs. In J. P. Dillard (Ed.), Seeking compliance: The production of interpersonal influence messages (pp. 87-104). Scottsdale, AZ: Gorsuch Scarisbrick.
O'Keefe, B. J. (1990b). Message design logic and the management of multiple goals. In K. Tracy (Ed.), Understanding face-to-face interacton: Issues linking goals and discourse (pp. 131-150). Hillsdale, NJ: Lawrence Erlbaum Associates.
O'Keefe, B. J. (in press a). Variation, adaptation, and functional explanation in the study of message design. In G. Philipsen & T. Albrecht (Eds.), Developing theories of communication. Albany, NY: State University of New York Press.
O'Keefe, B. J. (in press b). Developing and testing rational models of message design. Human Communication Research.
O'Keefe, B. J. & Delia, J. G. (1982). Impression formation and message production. In M. E. Roloff & C. R. Berger (Eds.), Social cognition and communication (pp. 33-72). Beverly Hills, CA: Sage.
O'Keefe, B. J., & Delia, J. G. (1985). Psychological and interactional dimensions of communicative development. In H. Giles & R. N. St. Clair (Eds.), Recent advances in language, communication, and social psychology (pp. 41-85). Hillsdale, NJ: Erlbaum.
O'Keefe, B. J., & Delia, J. G. (1988). Communicative tasks and communicative practices: The development of audience centered message production. In B. Rafoth & D. 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 intentions. Paper presented at the annual meeting of the Speech Communication Association, San Francisco, CA.
O'Keefe, B. J., & McCornack, S. A. (1987). Message design logic and message goal structure: Effects on perceptions of message quality in regulative communication situations. Human Communication Research, 14, 68-92.
O'Keefe, D. J. (1987, November). Message description. Paper presented at the annual meeting of the Speech Communication Association, Boston, MA.
O'Keefe, D. J. (1991). Persuasion: Theory and research. Newbury Park, CA: Sage.
Osgood, C. E., & Bock, J. K. (1977). Salience and sentencing. Some production principles. In S. Rosenberg (Ed.), Sentence production: Developments in research and theory (pp. 89-140). Hillsdale, NJ: Lawrence Erlbaum Associates.
Pavitt, C. (1991). An analysis of artificial intelligence based models for describing communicative choice. Communication Theory, 1, 204-224.
Perrig, W., & Kintsch, W. (1985). Propositional and situational representations of text. Journal of Memory and Language, 24, 503-518.
Pickering, A. (in press). Openness and closure: On the goals of scientific practice. In H. LeGrand (Ed.), Experimental enquiries: Australian studies in the history and philosophy of science. Dordrecht: Reidel.
Pickering, A. (1990). Knowledge, practice and mere construction. Social Studies of Science, 20, 682-729.
Pinker, S., & Prince, A. (1988). On language and connectionism: Analysis of a parallel distributed processing model of language acquisition. Cognition, 28, 73-193.
Regier, T. (1991). Quickprop.c [Computer program]. Berkeley, CA: University of California.
Ritter, H., & Kohonen, T. (1989). Self-organizing semantic maps. Biological Cybernetics, 61, 241-254.
Rogoff, B., & Lave, J. (1984). Everday cognition: Its development in social context. Cambridge, MA: Harvard University Press.
Rumelhart, D. E., Hinton, G., & McClelland, J. L. (1986). Distributed representations. In D. E. Rumelhart & J. L. McClelland (Eds.), Parallel distributed processing: Explorations in the microstructure of cognition, Vol. 1: Foundations (pp. 77-109). 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.
Rumelhart, D. E., & Norman, D. A. (1988). Representations in memory. In R. C. Atkinson, R. J. Hernstein, G. Lindzey, & R. D. Luce (Eds.), Steven's handbook of experimental psychology, Vol. 2: Learning and Cognition (pp. 511-587). New York: Wiley.
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.
Sacerdoti, E. (1974). Planning in a hierarchy of abstraction spaces. Artificial Intelligence, 5, 115-135.
Sacerdoti, E. (1977). A structure for plans and behavior. New York: Elsevier North-Holland.
Sacks, H., Schegloff, E., & Jefferson, G. (1974). A simplest systematics for the organization of turn-taking in conversation. Language, 50, 696-735.
SAS Institute, Inc. (1985). The VARCLUS procedure. In SAS user's guide: Statistics,

version 5 edition (pp. 801-816). Cary, NC: SAS INstitute, Inc.
Schank, R. C., & Abelson, R. P. (1977). Scripts, plans, goals, and understanding. Hillsdale, NJ: Lawrence Erlbaum.
Schlesinger, I. M. (1977). Production and comprehension of utterances. Hillsdale, NJ: Lawrence Erlbaum Associates.
Searle, J. R. (1969). Speech acts: An essay in the philosophy of language. Cambridge: Cambridge University Press.
Searle, J. R. (1975). Indirect speech acts. In Cole, P. & Morgan, J. (Eds.), Syntax and semantics, Vol. 3: Speech Acts (pp. 59- 82). New York: Academic 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.
Selman, R. L. (1980). The growth of interpersonal understanding: Developmental and clinical analyses. New York: Academic Press.
Shantz, C. U. (1975). The development of social cognition. In E. M. Hetherington (Ed.), Review of child development research, Vol. 5 (pp. 93-125). Chicago, IL: University of Chicago Press.
Sibun, P. (1990). The local organization of text. In Proceedings of the Fifth Internatonal 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.
Simon, H. A. (1977). Models of discovery. Dordrecht, Holland: Reidel.
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. Oxford: Basil Blackwell.
Suchman, L. A. (1987). Plans and situated actions: The problem of human-machine communication. Cambridge: Cambridge University Press.
Waldron, V. R. (1990). Constrained rationality: Situational influences on information acquisition plans and tactics. Communication Monographs, 57, 184-201.
Waterman, D. A., & Newell, A. (1971). Protocol analysis as a task for artificial intelligence. Artificial Intelligence, 2, 285-318.
Waterman, D. A., & Newell, A. (1972). Preliminary results with a system for automatic protocol analysis (C. I. P. Working Paper No. 211). Pittsburgh, PA: Carnegie Mellon University.
Waterman, D. A., & Newell, A. (1973). An interactive task-free version of an automatic protocol analysis system. In Proceedings of the Third IJCAI. Menlo Park, CA: Stanford Research Institute, 431-435.
Werner, H. (1957). The concept of development from a comparative and organismic point of view. In D. B. Harris (Ed.), The concept of development (pp. 125-146). Minneapolis, MN: University of Minnesota Press.
Wertsch, J. V. (1985). Vygotsky and the social formation of mind. Cambridge, MA: Harvard University Press.

APPENDIX A

INTERPERSONAL COMMUNICATION SURVEY

Name _____

We would like you to tell us how you would handle several different interpersonal communication situations. In this questionnaire you will find descriptions of several situations a person might find him/herself in. We would like you to read each description carefully and try to imagine yourself actually being in the situation. After you read the description, we will ask you to tell us, by writing your own words, what you would say to handle the situation.

For each situation, we would like you to write what you would say; do not describe the general action you would take--instead try to put in your own words what you would actually say to deal with the situation.

When you have finished reading these instructions, turn the page and begin.

GROUP PROJECT SITUATION

Imagine that you are enrolled this semester in a small group communication class in which a group project is assigned. It is important to you to get a good grade in the class, and your course grade will depend to a great extent on how well the group project turns out. The instructor simply assigned you to a group, and also assigned you to be the leader of your group. As group leader, your responsibilities will include meeting with the course instructor to assist in assigning grades for the project to your group members.

One group member, Ron, seldom makes it to group meetings on time and entirely skipped one meeting without even calling anyone to let the group know in advance. When Ron missed the meeting, two of the group members wanted to have the instructor remove Ron from the group, although another member persuaded the group to give him another chance. At the next meeting Ron arrived late but apologized for missing the previous meeting and mentioned something about family problems. Ron did volunteer to do all the background research on one important aspect of the group's topic, saying he had a special interest in that part of the project.

The group project is due next week. The group planned to put together the final draft of its report at a meeting scheduled for tomorrow afternoon. Ron calls you and says he doesn't have his library research done and can't get it finished before the meeting. He says he just needs more time.

What would you say to Ron? In the space below, write what you would say; DO NOT DESCRIBE THE GENERAL ACTION YOU WOULD TAKE--instead try to PUT IN YOUR OWN WORDS WHAT YOU WOULD ACTUALLY SAY to deal with this situation.

_____

_____

_____

_____

_____

_____

_____

(USE BACK OF PAGE IF NECESSARY)

BROKEN DATE SITUATION

Imagine you have a friend, Terry, and you and Terry have been close friends for three years. This last year Terry never seems to have time for you. Terry has cancelled a number of activities the two of you have planned because of meetings and other commitments. Your patience has been wearing thin lately: in the last week Terry has cancelled out on a lunch engagement and on a plan to go bowling with you.

You have plans to see a movie with Terry tonight at o'clock. You are really looking forward to seeing this movie and to having a good time with your friend. At three o'clock Terry calls you to cancel out. Terry says, "I'm really sorry but there's a meeting at : I really have to attend. I would have called sooner but I just found out about it. Can we go to the movie tomorrow night?"

What would you say to Terry if you found yourself in this situation? In the space below, write what you would say; do not describe the general action you would take--instead try to put in your own words what you would actually say to deal with this situation.

_____

_____

_____

_____

_____

_____

_____

_____

_____

_____

(USE BACK OF PAGE IF NECESSARY)

Now we would like to know what you were thinking as you decided what to say to Terry. The following pages list a number of different kinds of thoughts that peole have told us they had as they responded to this situation. Of course, different people have different thoughts about a situation like this one: some of the thoughts on this list will be ones that occurred to you, but others will not be similar to your own thoughts.

We would like you to read throught this list and mark an X in the space next to any thought that you recognize as one that occurred to you whaen you read and responded to the situation involving Terry. If a thougt seems very similar to one of yours, but just worded differently, then mark it with an X even though it may be a little different from your own. But don't just select a thought from his list because it is somewhat similar to your own--only pick thoughts from he list that you fel are the thoughts that occurred to you as you imagined the situation and decided what to say.

Some of the thoughts on the list that didn't occur to you when you first read the situation may strike you now as being very true. Do not select thoughts from the list just because you now see they re true. We are interested in knowing which thoughts you were thinking as you composed your message, not which ideas you think are true, so select only thoughts you recognize as your own natural response to the situation.

When you have finished reading these instructions, turn the page and begin.

Which of the following feelings did you have?

_____ I was looking forward to seeing Terry tonight.
_____ Terry's actions hurt my feelings.
_____ Terry irritates me.
_____ It's okay to cancel once in a while, but not repeatedly.
_____ I'm tired of being patient with Terry.
_____ I'm upset with Terry.
_____ I'm sick of being let down.
_____ I'm tired of Terry's behavior.
_____ I feel disappointed.

Which of the following thoughts about the relationship did you have?

_____ I'm busy too, but I don't let that interfere with seeing Terry.
_____ The friendship is important to me.
_____ Talking with Terry about the friendship is important.
_____ I value Terry's friendship.
_____ Terry is a close friend.
_____ Terry is a good friend.
_____ Terry isn't behaving like a good friend.
_____ If I criticize Terry it will hurt our friendship.
_____ This isn't the only friendship I have.

Which of the following thoughts about Terry's previous actions did you have?

_____ Whenever we have plans something else always comes up.
_____ I understand why Terry had to cancel our date.
_____ Terry repeatedly cancels appointments with me.
_____ Terry want to do the movie tonight but can't because of a movie.
_____ Terry's telling the truth.
_____ Terry doesn't seem to have time for our friendship.

Which of the following thouhgts about Terry's traits did you have?

_____ Terry's irresponsible.
_____ Terry was considerate to call me.
_____ Terry acts responsibly about attending meetings.
_____ Terry is selfish.
_____ Terry is inconsiderate.
_____ Terry is disorganized.
_____ Terry's a liar.
_____ Terry is ambitious.
_____ Terry is a happy person.
_____ Terry is insecure.
_____ Terry is a leader.

_____ Terry is active.
_____ Terry is highly motivated.
_____ Terry is insensitive to the feelings of others.
_____ Terry is overcommitted.
_____ Terry is busy.
_____ Terry's guilty.
_____ Terry knows too many people.
_____ Terry tries to be too involved.
_____ Terry is self-centered.
_____ Terry thinks people should work around his/her schedule.
_____ Terry is rude.
_____ Terry is an overachiever.

Which of the following thoughts about Terry's motives and feelings did you have?

_____ I don't know what Terry's motives are.
_____ Terry's using activities as an excuse not to see me.
_____ Terry has a problem with me that he/she is afraid to bring up.
_____ Terry feels our friendship is stale.
_____ Terry doesn't consider the needs/feelings of others to be important.
_____ Terry doesn't place a high priority on our friendship.
_____ Terry thinks it is important to keep commitments.
_____ Terry has personal problems and obligations I don't know about.
_____ Terry doesn't realize that cancelling apointments can hurt others.
_____ Terry cares about me and my feelings.
_____ The friendship is important to Terry.
_____ Terry would like to see me tomorrow night.
_____ If Terry knew there was a problem between us, he/she would be worried.
_____ Terry wants to break off our friendship.
_____ Tery would like it if the situation were cleared up.
_____ Terry feels bad about cancelling the date.

Which of the following thoughts about how Terry should behave did you have?

_____ Terry takes me for granted.
_____ Terry's avoiding me.
_____ Terry doesn't respect me.
_____ If I refused to see Terry, Terry would miss me.
_____ Terry expects me to be understanding.
_____ Terry thinks I want to do things with him/her.
_____ Terry wants me to understand what is happening in this situation.
_____ Terry would respect me for being firm.
_____ If Terry knows I feel hurt, it will make me look bad.
_____ If Terry has a problem with me, he/she should say so.

_____ Terry ought to have time for me.
_____ Terry should apologize for being a bad friend.
_____ Terry should put more effort into our friendship.
_____ Terry ought to explain the repeated cancellatons.
_____ Terry ought to take responsibility for our friendship.
_____ Terry should promise not to cancel again.
_____ Terry should be put on the defensive.
_____ Terry should organize time better.
_____ Terry should keep appointments.
_____ Terry should make fewer commitments to others.
_____ Terry should cancel other commitments rather than dates with me.
_____ Terry should stop acting interested in seeing me and then cancelling.
_____ Terry should skip the meeting tonight.
_____ Terry should call me back when he/she has time for a long talk.
_____ Terry should be the one to make plans in the future.

Which of the follwoing thoughts about your own actions did you have?

_____ I should stop seeing Terry for a while.
_____ I should let Terry know how I feel about the cancellation.
_____ I don't want Terry to know that I feel hurt.
_____ I want to do to Terry what he/she did to me.
_____ I should be supportive to Terry.
_____ I should be understanding about Terry's problems.
_____ I should refuse to go with Terry tomorrow.
_____ I should give Terry another chance.
_____ I should be firm in dealing with Terry.

Which of the following thoughts about hoe=w Terry should feel or think did you have?

_____ Terry should be grateful to me for being understanding.
_____ Terry should know that I understand about the cancellation.
_____ Terry should realize what he/she has been doing to me.
_____ I want Terry to feel bad.
_____ Terry ought to see that cancelling appointments is bad behavior.
_____ Terry needs to know that I feel hurt.
_____ Terry needs to know that I am a little angry.
_____ Terry needs to understand the consequences of cancelling appointments.
_____ I want Terry to feel guilty.

Which of the following thoughts about future plans involving Terry did you have?

_____ Terry will try to keep the appointment tomorrow night.
_____ I could go to the movie with someone else.
_____ Terry and I could do something else and skip the movie.

_____ If I agree to go tomorrow, Terry will probably cancel that too.
_____ I can make it tomorrow.
_____ I should not make dates with Terry in the future.
_____ Our friendship needs to be put back on track.
_____ I want to spend time with Terry.
_____ I realy want to see this movie.
_____ We shouldn't make definite plans again until Terry has more free time.
_____ I want Terry to go out with me.
_____ We need to get things out in the open.
_____ I want more control in this relationship.
_____ It would be better if Terry weren't so sure of my friendship.
_____ We should get together at a time of Terry's choosing.

APPENDIX B

Rules for Coding the Presence/Absence of Thoughts

The items on the Thought Checklist were condensed from a much larger set of thoughts collected in a pilot study (O'Keefe & Lambert, 1989). In designing the Checklist, standard forms of expression were chosen to stand for whole sets of synonymous thoughts. In this way, the Checklist could capture the largest range of thoughts in the smallest number of choices. Thus in determining whether a protocol-elicited thought was or was not on the Checklist, a standard of synonymy somewhat less stringent than literal matching had to be devised. Protocol thoughts were coded as either present on or absent from the Thought Checklist according to the following guidelines.

1. Synonymous thoughts. Thoughts synonymous with Checklist thoughts are coded as present.

Examples:
"Terry pisses me off" counts as "Terry irritates me"
"I'm sick and tired of Terry's bullshit" counts as "I'm tired of Terry's behavior."
"It's not that Terry doesn't care about me" counts as "Terry cares about me."

2. Questions count as propositions. No questions apear on the Thought Checklist, so questions from the protocols are transformed into propositions.

Examples:
"Why doesn't Terry have time for me?" counts as "Terry doesn't have time for the friendship."
"Why are you doing this Terry?" counts as "I don't know what Terry's motives are."

3. Quotes. Direct quotes are compared to standardized expressions.

Examples:
"Call me when you want to go" counts as "We should get together at a time of Terry's choosing."
"I'd tell Terry he had a little too much on his plate" counts as "Terry is overcommitted."

5. Requests for accounts. Requests that Terry account somehow for his prior behavior are

generally coded as present.

Examples:
"I'd ask Terry why he'd been blowing me off" or "Why do you always cancel our dates?" counts as "Terry should explain these repeated cancellations."

6. Thought complexes. When several thoughts on te protocol are equivalent to one thought on the Checklist, each thought in the complex is coded as present.

Examples:
"It was okay / Terry cancelled a couple of times / Then it got bad" counts as "It's okay to cancel but not repeatedly."

APPENDIX C

Rules for Coding Protocol-Specific Thoughts

The following coding categories were derived from protocol data exclusively. These categories represent types of content that originate as much out of the social setting of the protocol task as out of the Broken Date scenario itself.

1. Construing replications of events. These elements refer to the subject's previous experiences, experiences that resemble events in the Broken Date scenario.

Examples:

This has happened to me before with other friends.
I know someone named Terry.
I've been in this situation.
This reminds me of my roommate.
I immediately associated that with a true life story.
Just like me and Sara.
I was recently having something similar with my friends.

2. Questions or comments directed at experimenter. These elements are questions or comments about the experimental procedure that were directed at the experimenter.

Examples:

Is Terry a boy or a girl?
I'll assume Terry's a girl, okay?
Is Terry a friend from school or from some activity I'm involved in?
Do I keep reading here?
That [the scenario] is pretty descriptive.
How long have I planned to see the movie?
Is Terry an organized person?
Oh, this is one of those situations.
I'm assuming this is a face-to-face confrontation.
This page seems to contradict the last one.
Am I supposed to write and think at the same time?

3. Descriptions of subject's own mental state. These elements are subject's characterizations of their own mental states.

Examples:

That's all I can remember.
I can remember what I thought but I can't remember what the situation was.
I don't know.
I changed my mind.
That's kinda the gist I got from that last page.
My thoughts were more negative than what I said.
My response was a little more understanding than my thoughts.

4. Non-specific subjective reactions. These elements include generalized subjective reactions to the Broken Date scenario as it unfolds.

Examples:

That is bad.
I think that's great.
That's not good.
Humm, this is just not good.
That's good I guess.
That's not right.
This is hard.
That's nice.

5. Asides. These elements were uttered as humorous asides to the experimenter.

Examples:

I hate bowling.
Bowling is my favorite sport.
Oh the typical "do lunch" thing and never show up and never call which is so typical of U.of I.

6. Preambles. These elements preceded the reporting of other substantive thoughts.

Examples:

I'd kinda handle this in two ways.
I'm not sure what I'd say to Terry.
I would tell Terry this.
What was I saying?

I'll just make it up.
I'd do one of two things.

7. Inaudible elements. These elements were not clear enough to be accurately transcribed.

APPENDIX D

Message Elements Sorted by Message Frame

Frame #1: You have broken our dates a lot lately.

You have missed several dates.
I feel like you are blowing me off.
That is why you are blowing me off.
You are just having trouble keeping commitments.
Why do you keep ditching out on me?
This is the third time this week you have canceled out on me.
This has come on so suddenly.
Terry you have done this so many times lately.
You have canceled out on me too many times this past year.
Why do you keep breaking our plans.
You never go when you say that you want to.
You have been canceling all of our engagements.
You have been backing out on everything lately.
Every time we make plans you seem to have something better to do.
Something suddenly comes up.
Terry why do you keep doing this to me.
All I have gotten so far this year are excuses.
Terry why am I not surprised that you canceled.
Every time we have plans you find something more important to do.
This is not the first time something like this has happened recently.
Listen Terry you have really been kind of blowing me off lately.
Every time we have something planned you have something more important to do.
Whether you have the meeting or not you have been blowing me off big time lately.
Every time we plan something you cancel out.
You have made promises before and canceled out.
Can you explain to me why you always break our plans.
You always break our date at the last minute.
Why do you keep canceling our dates.
This is the third time you have canceled.
Do you realize we have not carried out our plans to do something in weeks.
Every time we make plans you cancel out.
Now you have ruined my plans again.
You have been canceling out on me a lot lately.
Every time we have plans you cancel.
You always cancel your plans.
Recently you have canceled and pulled out of every plan we have had to do something together.
You realize you have canceled the last three engagements we have made.

Terry you have been blowing me off lately.
You keep canceling out.
This past week you have missed lunch bowling and now the movie.
I really think you are neglecting me.
You cancel all the time.
Terry do you realize how many times you have canceled our lunch and other engagements.
Terry you know this happens all the time.
You never keep plans that either of us makes.
Why do you cancel so much.
All our other plans seem to be canceled so often.
You have been constantly canceling activities with me.
You never follow through.
You do this to me so often.
You have blown me off numerous times.
This is the third date this week you have broken.
You continue to back out.
Look dude you have been skipping out on our plans all the time.
Why is it I always seem to be blown off.
It seems as if you have been really blowing me off lately.
I do not like it that you always call at the last minute and cancel plans.
I can not believe you canceled again.
You keep blowing me off.
You have blown me off every time we have tried to go out these past weeks.
Every time I am looking forward to something you cancel out.
Why do you keep canceling out on me.
Terry this is becoming a habit with you.
Terry why do you keep canceling on me.
I can not believe you are blowing me off again.
You are the one who always cancels.
This canceling bit has happened a few too many times.
In the past weeks you have been canceling our dates frequently.
Terry this is the third time in two weeks that you have blown me off.
Well that should not surprise me that you have canceled again.
You have broken so many dates that we were supposed to go on.
You have blown me off three times this week.
This always happens.
Lately all you have been doing is backing out of them.
Terry you have really been blowing me off lately.
That is the third time this week.
Terry why is it you always seem to back out on our plans for these important meetings.
You make plans with me then end up canceling out because of other things.
Terry you have cancelled so many things we have planned to do together.
You have broken so many plans with me.
Over the last year you have been constantly blowing me off for your damn meetings and other shit.

I mean it has been a full year already of you making and breaking our plans.
I guess not because you keep backing out of our plans.
Terry this is the third time this week you have canceled our plans.
You keep cancelling out on me.
You have been blowing me off big time.
You are always cancelling out on me.
It seems as though whenever we finally make plans you cancel on me.
Why do you continue to plan activities with me only to break them off at the last minute.
You have already done that twice this week.
Well Terry old buddy old pal of mine here is yet another time you cheesed out on our plans.
Now you are always canceling on me.
The last ten times you have backed out on me I took it with a grain of salt and have not really said anything about it.
You just cancel all of our engagements anyway.
You know you are going to back out.
You know Terry it always seems like it is something.
Somehow our plans always get cancelled.
Somehow our plans always fall apart.
You have blown me off for two weeks straight.
In the past weeks you have cancelled everything we have had planned.
Then you are backing out.
Every time we make plans to go out you cancel.
You keep blowing me off.
This is the third time this week you have canceled our plans.
I want to point out to you how frequently you have been canceling plans with me.
Do you realize you have canceled plans together for the last few weeks.
Now you cancel on the movie.
This is the third time in a week that you have done this.
It seems to me that you are blowing me off once again.
Every time we make plans you cancel.
You have done this to me three times now.
This has been happening a lot lately.
Every time we make plans together you break them on me.
You have broken several lately.
Terry this is the third time this week you have canceled our plans.
You have canceled many more before that.
You have been blowing me off a lot lately.

Frame #2: I'm tired of your behavior.

I am just tired of being stood up and canceled out by you.
I am really getting tired of you postponing all of our plans.
I am very fed up with the cancelations you have been making lately.
I would like you to know I am really getting tired of your backing out on all of our plans.

I am getting kind of sick of you always canceling out.
I am sick of you blowing me off.
I am really sick of being blown off.
Terry I am sick of the way you have been treating me lately.
I am getting sick of your bs.
Terry I am getting sick and tired of you ruining all of our plans all the time.
Listen Terry I am really getting fed up with your breaking plans.
Terry I am sick and tired of you canceling on me.
I am sick of it.
I do not need this.
I am sick of your excuses.
I am sick and tired of your meetings.
I am sick of working around your schedule.
I do not need that.
Terry I am tired of having you cancel out all the time.
I am sick of being blown off.
I am really getting tired of this Terry.
I am really hurt.
I am tired of being disappointed.
Terry I am really tired of your games.
I am really sick of making plans with you and having them broken all the time.
I am really tired of you canceling dates with me.
I am sick of this.
I am getting tired of your constantly breaking our date.
I am tired of getting canceled on.
Enough is enough.
I can not handle much more of this.
I just want to say I am sick and tired of being your little puppy.
I am tired of making tentative plans with you and having you cancel them.
I am sick of it.
I am sick and tired of being used as a friend when you have nothing better.
I am really getting a little bit tired of getting blown off for one of your numerous activities.
I am tired of you blowing me off.
I am tired of you saying you have too much to do.
I am beginning to get sick and tired of you always putting me and our friendship last on your list.
I am tired of getting ready to go out and then being stood up.
I am tired of being happy one minute and let down the next.
I am getting really sick of you canceling out on me for the past year.
I am sick of it.
I am sick and tired of hanging by a string.
I am tired of all this bs you have been giving me.
I am tired of you blowing me off.
I am tired of trying.
I am really tired of you breaking plans with me.

I am tired of hearing I have a date at eight o 'clock.
I have to do this.
I have to do that.
It is tiring.
I am just tired of it.
I am getting kind of tired of it.
I can not handle this.
I am sick and tired of putting my life on hold for you.
I am tired of constantly blowing off plans.
I am tired of having plans with you and then getting canceled on.
I am getting sick of being blown off.
I am sick of the way things have been going.
Terry I am damn sick and tired of your rudeness.
I am kind of tired of reserving time for you only to have you blow me off.
I am getting sick of being blown off by you.
I am tired of being the only one putting forth any effort.
I am tired of it.

Frame #3: Your behavior is bothering me.

I feel hurt that you have canceled all of our previous activities.
I am kind of mad.
Terry your behavior is really upsetting me.
I am really hurt by the cancelations.
I get so hurt.
I feel really left out.
Terry you are a jerk.
I do not love being canceled on.
I want you to understand that it is important to me that we go tomorrow.
I just wanted to tell you that I really do not appreciate your behavior at all.
It hurts.
I am upset with you Terry.
You are really pissing me off.
I feel almost as though my hopes are being crushed every time you cancel.
I am really hurt about what you have been doing.
I think you are inconsiderate.
This canceling bit is really annoying.
The situation was upsetting me.
I am upset that you keep canceling our engagements.
I want you to know I sometimes feel bad when you cancel.
To be honest with you I am hurt every time you cancel.
I am continually feeling neglected.
It really hurts when these last minute things come up.
It is pissing me off.
Terry you are really making me mad lately.

It hurts that every time we make plans to go out and have a good time you cancel.
I would not mind if you only canceled once in a while.
Terry it has really been bothering me the way you have been neglecting me recently.
It is a real disappointment for me when you cancel.
Terry something has been bothering me for a long time now.
You are hurting my feelings.
It has been really frustrating.
I hope you realize that it does get a bit annoying to me when someone is constantly breaking dates or plans with me.
I do not care.
I do not know if you realized that this is something that is upsetting me.
I feel hurt by the way you have been treating me.
Terry I am really disappointed that we will not go to the movie together tonight.
I think you have no idea what the word commitment means.
You are just making my life miserable.
I am angry that you canceled on such short notice.
Right now I feel like I have been left in the dust.
I hate making plans and then having them fall through.
It is getting pretty irritating.
I feel hurt that you keep canceling on me.
I feel hurt and rejected when you cancel our plans every time something comes up.
It is hard to take being blown off all of the time.
I am really mad that you never seem to have time for me anymore.
All of these last minute cancellations are becoming a real inconvenience.
Lately I feel left out.
I feel bad when you cancel on me.
This not only hurts me.
It makes me upset.
I hate not to take advantage of my free time.
I do not appreciate being blown off like the way you have been doing to me in the past year.
I am too frustrated with what has happened already.
Please do not continue treating me as you have.
It really bothers me.
Terry I really wish you would not cancel at the last minute.
It really gets irritating.
I am disappointed in you.
It has upset me that we have not seen each other.
I feel like I am being dumped all the time.
I always feel like a jerk a loser.

Frame #4: Your behavior is unacceptable.

You are kind of irresponsible in terms of scheduling.
I know it is selfish.

You are unsensitive and uncaring.
It is unfair.
It is not fair.
This meeting stuff all the time is getting to be ridiculous.
You really piss me off.
You know you really have scheduling problems.
This is crazy.
Three cancelations in one week is ridiculous.
Give me a break.
Terry I have had it.
This lack of consideration is really rude.
You know all your other friends would have never taken this attitude from you.
You just do not treat someone like you have been treating me lately.
Oh come on Terry.
It is not fair for you to make dates you can not keep.
What you have been doing is rude.
It is not cool.
You have done nothing but jerk me around lately.
You have been acting really weird lately.
Listen Terry dude you have got to stop pulling this bullshit with me.
It is bullshit.
All this canceling of dates has got to stop.
I do not think it is fair of you to keep blowing me off like this.
Terry I can not believe you are doing this to me.
This is the last straw.
Anyway this busy excuse is overused.
You are screwing up my scheduling process.
Terry I want you to know you are not being fair to me.
It is rude unfair and inconsiderate.
This is getting ridiculous.
I can not deal with cancelation after cancelation anymore.
It sucks.
This has gotten out of hand.
It is not considerate to me to keep canceling.

Frame #5: What is going on?

What is up with you?
I do not know what your problem is.
Your other commitments cannot take up all of your time.
Terry I do not understand.
Which is it?
Do you want me to stop begging you?
What is going on lately anyway?

What is the deal?
Do you just not want to go?
I do not know what the problem is.
What has happened?
It is hard to understand your recent business.
Please let me know what is happening.
What are you doing that is so important that you can not do anything with me anymore?
What is the deal?
What is happening?
What the hell is going on?
What is the deal?
What is the matter?
I just do not understand what is going on here.
You are not making any sense to me.
I could understand if something popped up once in a while.
What is up?
It is all right if you have to cancel a few of our engagements.
What is going on?
What is your problem?
What is up with our relationship?
Tell me what is going on.
I want to know what is going on.
What is going on Terry?
How is your homework?
I just do not understand why you keep doing this to me.
What is going on in your life that makes you so scatterbrained that you can not even plan to meet with me without eventually cancelling?
Terry what the fuck is going on?
Why the hell do you make plans if you know you have other commitments?
Tell me what is going on.
So how about telling me what is going on.
I just do not understand what is going on now.
Do you still want to be friends with me?
Why are you being a dick to me?
What am I not cool anymore?
Is it because we are just friends?
Just tell me what is going on.
What is the story?
Hey man what has been up with you lately?
Why have you been acting so bogus to me lately?
I just do not understand what is going on.
What do you want me to do plan something to go with each other at three in the morning or at some other awkward time?
Why do you even bother to arrange to do anything?
Do you feel at all bad about canceling our plans all of the time?

Are you really blowing me off?
What you do this all the time?
What is so important this time?
What is the deal?
What is going on?
I do not understand you.
Why should I be understanding and accept your proposal to go tomorrow night?
Terry I do not know what is wrong or going on here.
Terry I do not understand why you keep doing this.
What is the problem?
What is going on with us?

Frame #6: Is there a problem?

Tell me what is wrong.
Have I done something to make you want to avoid me?
Is there something wrong?
Are you mad at me?
I am not boring or something am I?
Is anything wrong?
Is there something wrong?
Are you mad at me about something?
Is there something wrong?
Is something wrong?
Is this the case?
Did I do something?
Is there something wrong that is bothering you?
Is there anything I can do?
Is there another guy?
Are you unhappy with our times together?
Are you mad at me?
Is something wrong?
Have I done something to make you mad?
Why can not we just get together?
Are there any problems I should know about?
Is it something I did or said?
Have I done something wrong?
Is it something I did?
Is there something other than the meeting keeping you from going out?
Is that it?
Have I done something to upset you?
Is there something wrong?
Are you mad at me for some reason?
Is there something wrong?

If there is some problem please tell me?
Is it something I did or said?
Are you pissed at me?
Did I do something to make you mad?
I do not think I have?
Terry is there something wrong?
Terry I just want to know if something is wrong?
Did I do something to upset you?
If you are mad at me please let me know?
Terry are you trying to tell me something?
Is there something you need to talk about?
Is it me?
Is it something I have done?
I do not even think I did anything?

Frame #7: You've been busy.

Why are you so busy when we make these plans.
You are busy.
You have not had much time lately to see me.
Terry are you too busy for me or what.
I realize that you are busy.
Lately you have been really busy.
I understand you are busy.
I understand you are very busy Terry.
Are you really busy.
I am worried that your commitments will be too much to handle during school.
Maybe you are simply busy.
It is like you are too busy.
I understand you have been busy lately.
I know you are busy and stuff.
You have been so busy lately.
You have all these other commitments to take care of before you will do anything with me.
I guess you are just too involved in activities.
I know you are busy.
I understand you are really busy.
I know you are real busy.
Maybe you are getting in over your head.
Terry I know that you are involved in a lot of things.
You are busy.
I know that you are involved with many things that are important to you.
If you do not have time for friends maybe you are too busy.
I understand that you are a very busy person.
It seems that you have been pretty busy lately.
I understand that you are very active now.

You seem to have a really busy schedule and not have time for outside activities.
I know you are really involved.
I know that this year you are very busy.
I know you are busy.
It is just that lately you always have something else.
I know you are real busy.
It seems like you have been awfully busy lately.
I realize that you are really busy this year with meetings and other commitments.
You are always busy.
Terry I know you have been busy with classes and your extracurricular activities.
I know you are really busy this year.
I understand you are busy.
I understand you are busy.
Listen Terry I understand that you are very busy.
Terry I know you have been busy lately.
I understand you are busy.
I know you have other things you have to do.
I know you have a busy schedule.
I know you have other responsibilities.
I know you have a tight schedule.
You seem really busy these days.
Listen I know you are busy with all sorts of other things.
I understand you are busy.

Frame #8: Forget it, Terry.

Let us just forget about it okay.
Go to hell.
Forget this bullshit.
Go screw over someone else.
Do not call me I will call you.
Do not bother making the plans.
No thanks for tomorrow night and for a while for that matter.
Till then goodbye.
Terry look let us just forget it.
Forget the movie.
It would be better if we did not go to the movies tomorrow night.
Forget it.
I think it would be better if we forgot about the movie altogether.
Hit the road Terry.
Forget about the movie for tomorrow night.
Fuck it.
Just forget it Terry.
Kiss off.

I suggest you get a friend who will put up with your crap.
Do not bother calling me ever again.
Maybe we should just skip it all together.
Listen Terry forget the movie.
Just forget it okay.
I do not need any bullshit.
I do not care about your meeting.
Why should I care about your stupid meeting.
Just forget about doing anything anymore.
I have had it.
Right now I really do not care.
You know what Terry let us just forget it.

Frame #9: You don't have time for me.

All you have had time for are meetings meetings meetings.
You have no time for me anymore.
Why can not you find time for me.
I can not understand why you can not find any time to spend with me.
It does not seem like you have time for me anymore.
You never make time for me.
I feel as if you do not have time for me.
Why do not you have time for me anymore.
That is all you seem to have time for.
Terry it seems that you never have time to do anything with me anymore.
It seems you never have any time for our friendship.
It seems like you never have time to do things together anymore.
I wish you would know what is going on with your life that you no longer have time for me.
It seems you have no time for me lately.
You never seem to have time to socialize anymore.
Terry you have not had much time for me lately.

Frame #10: I suspect that you are avoiding me.

You are not admitting to me that you do not want to spend time with me.
It seems as of late you have been trying to avoid me.
I guess I want to know why you do not want to spend time with me this year.
You never seem to want to do anything with me lately.
I do not understand why you do not want to do things with me anymore.
I feel as if you do not want to be around me.
I feel as if you do not enjoy my company anymore.
I am beginning to see a pattern here.
Do you still want to be friends with me.

You do not like me any more do you.
Do not you like me anymore.
I was wondering why have you been avoiding me.
I feel as if you do not want to do anything with me anymore.
This last year you seem to avoid me by blowing me off all the time.
I get the feeling that you are canceling because you do not want to see me.
Why have you been avoiding me.
If we are such good friends why do not you want to do anything with me.
Lately it seems that you are trying to avoid me.
I find it a little fishy that every time I have a date with you you seem to have something come up.
It seems obvious that you really do not care to see me.
I get the impression you do not want to go out with me.

Frame #11: We haven't been together for a while.

We have not done anything in a good long while.
I do not even remember the last time we went out together do you.
It seems we never see each other.
I have not spent any time with you.
We never get together.
We never do anything together anymore.
We have not really been together now for a year.
I have not seen you in ages.
We have not gone out in such a long time.
I mean I never get to see you.
We have not spent a lot of time together lately.

Frame #12: We have done so much together.

We have done a lot together.
We used to do so much together.
I always thought we were having fun.
We used to spend a great deal of time together.
We used to go to a lot of places together.
We have done everything together.
We have had some really good times.
We always used to hang out.
Remember them.

Frame #13: I want to get together.

I hope to see you tomorrow.
I do enjoy spending time with you.
We really should get together.
I would kind of like to do something together and get away from school and all the other stuff.
I would definitely like to go tomorrow.
I was really looking forward to going out with you.
I really enjoy being with you.
Can you please make sure we can get together sometime this week.
I would just like to get to see you that is all.
I want to spend more time with you.
I really want to get together.
I really want to see you.
I hope we can get together tomorrow.
Terry I really wanted to spend time with you tonight.
I look forward to doing things with you.
You are great to be with.
We will try to put something together.
I hope we can spend more time together in the future.
I was really looking forward to spending some time with you.
I would really like to go with you.
I wish we could spend more time together.
I was really looking forward to the movie and seeing you.
I would really like to spend some quality time with you.
I really wanted to go with you to that movie tonight.
We have to go out soon and have some fun.
I was really looking forward to spending time with you.
I hope we can plan something to do sometime.
Let us try to get together sometime soon.
Terry I really wanted to get together with you.
I want to get together.
I enjoy doing things with you.
I really enjoy spending time together.
I love spending time with you.
I really want to see you.
I would like it if we spent more time together.
I really want to spend time with you.
Terry I was looking forward to seeing you.
I want to be able to see you sometimes.
I was really looking forward to going out.
I was really looking forward to seeing you and going to the movie.
I was really looking forward to seeing you.
I am looking forward to it.

I just wish we could spend time together like the old days.
I really want to see you.
I like to do stuff with you.
I do want to get together soon.

Frame #14: I miss you.

I really miss you.
I really miss doing things together.
I miss our silly jokes and long talks.
I miss spending time with you.
I miss all the fun we have going out together.
I am a lonely person with your friendship and guidance.
I miss doing things with you.
Boy I miss all the good times and laughs we had.
I am jealous of all of the organizations you are involved in because they get you.
I really miss you Terry.
I miss you.

Frame #15: I want to remain your friend.

I do not want that to happen.
I hope I still am one of your friends.
I want to stay friends with you.
I hate to see our friendship dissolve.
I would hate to lose out on all the good times we can have.
I do not want you to be mad.
I would not want our friendship to end.
I do not want us to grow apart.
I hate to see our friendship end this way.
I really do not want that to happen.
I feel it would be crazy to lose that friendship because of your meetings.
I do not want to lose our friendship.
I hope you do not either.
I really like being your friend.
I want to stay friends with you.
I am not willing to throw our friendship away.

Frame #16: We have been good friends for some time.

We used to be such close friends.
You are my friend.

I am a good friend.
I thought we were good friends who could take time out for each other.
Terry you know we are good friends.
We are friends.
I thought we were friends.
We have been friends for three years.
You are a great friend.
We have been close friends for a long time.
I have been your friend for three years.
You are a close friend.
I realize our friendship is strong enough that if we do not see each other we will still be close.
We used to be close friends.
I am your friend.
Terry listen I thought I was your friend.
We used to be such close friends.
We have been friends for three years.
You are a very close friend.
Terry you are supposed to be a close friend.
You know I am your best friend.
You are really a good friend.
Terry we used to be such good friends.
We used to be great friends.
We have been friends for a long time.
I still consider you a very good friend.
I thought we were friends.
We have been close friends for three years now.
We have been best friends for years.
We have been friends for a long time.
I could always count on you.
I have really enjoyed running around with you for the past three years.
Terry we have been friends for a long time now.
Terry I know we have been friends for a long time.
We have been friends for so long.
We used to be boys.
I thought we were good friends.
We have been friends for three years.
I thought we were friends.
I mean we have always been with each other.
We have been buddies for a while.
I know that we have been friends for three years.
I thought that we were supposed to be friends.
I think of you as a really close friend.
I really thought we were good friends.
Terry we have been friends for three years now.
I thought we were close friends.

Listen Terry for the last three years we have been best of friends.
After all I am one of your best friends.
I consider you a good friend.
We will still be friends.
We have been friends for three years now.
We have been friends for a while.
We have been friends for a long time.
Sure you are my friend.

Frame #17: You have not been a good friend lately.

You are not a real friend.
I feel like you are not putting any effort into our friendship.
Lately you have not been a good friend.
I do not even feel like one of your friends.
I feel like you do not seem to have much interest in our relationship these days.
I feel as if you do not care.
I feel that recently our friendship means nothing to you.
I am sorry our friendship does not mean anything to you.
Maybe I am just not your friend.
Now you are just not there for me.
I deserve better than how I am being treated at this point in our relationship.
You do not care about my feelings.
Lately you have treated me like dirt.

Frame #18: Our friendship is important to me.

I value our friendship.
I am counting on your friendship.
I really want our friendship.
Your friendship is really important to me.
I need this friendship.
I value that friendship.
Is not our friendship worth something.
I guess our friendship means more to me than you.
I value our friendship.
Our friendship is to valuable to me.
Friendship is important.

Frame #19: I care for you.

I will always be there for you.

I really care about you.
I love being with you.
I want to be there to support you and your activities.
I like you a lot.
You will always be my friend.
I care for you a lot.
I am the one who is concerned about it so much.

Frame #20: Our friendship is suffering.

Something is wrong with our friendship.
I feel as if our friendship is drifting.
I feel like our friendship is on hold.
Lately it seems like our friendship is getting weaker.
Something has changed between us.
Our friendship is not important any more.
It seems we are drifting apart.
I think our relationship needs to be mended.
I thought we were better friends than we obviously are.
I guess ours is not a mutual relationship.
I feel as if our friendship is slipping.
We are growing apart.
It does not even seem like we are friends anymore.
It is more like we are acquaintances.
I feel like since I do not see you that often we are not as close as before.
Due to commitments that may be beyond your control our friendship has suffered.
I feel like we are growing apart.
I feel our friendship is falling apart.
I do not see it that way anymore.
It just seems as though we are drifting apart.
It does not seem like it.
It is hard to be close to you if you are always close to canceling our plans.
Because we have been growing apart I do not have the same feelings for you anymore.
You are letting our friendship slip away.

Frame #21: Our friendship may end.

I do not know what is going to happen to our friendship.
Our friendship is not going to stay as it is now.
It has to change.
If you do not tell me what the problem is our friendship is over.
Then I would never speak to you.
One of these days you will be really surprised to find that I will not be here for you to always

fall back on.
Are you attempting in any way to end our relationship.
If you do not want to be friends then keep on acting the way you are.
If our friendship has to end it has to end.
Unless things change we will not stay best friends.
I am going to reevaluate our friendship.
If your behavior does not stop I am going to have to start making other plans.
If you want this friendship to be one sided you are going to be the one.
You are going to lose a good friend if you do not put more into this friendship.
You are going to lose a good friend.
If this shit is not resolved soon you are going to end up one less friend.
We will not become enemies just because of all your time constraints unless you continue to make me plan my time for nothing.
Eventually I am going to run out of patience.

Frame #22: Friendship has certain requirements.

You have to be fair to your friends.
Is that really what friendship is.
Friends do not treat friends as if they were dog shit.
To have good friends you have to be a good friend.
Friendship takes two people.
I can not make our friendship work alone.
I can not support you unless you support me.
A relationship needs commitment and give and take between partners.
A friendship is not just something to throw on the back burner until you have time for it.
I mean what are friends for if you do not have time to spend with them.
Friendship is a give and take situation.
I thought friends were supposed to be there for other friends.
It is a two way street.
Friends do not make plans and promises and then constantly break them.
Friends do not just do this to each other.
A one sided friendship is not a true one.
Friends are supposed to be forgiving.
Friends are also supposed to care.

Frame #23: We need to talk.

I really want to talk about this.
If we are growing apart we have to decide what the best move for each of us is.
We need to talk.
We have lots of catching up to do.
We need to talk.

Terry I have been meaning to have a talk with you.
I think we should talk.
Terry I think we need to talk and get some things straightened out.
Terry we really have to talk.
I wish you would talk to me about the problem.
I think we have to talk.
I have to talk with you.
It is about time we discussed the problem that has developed in our friendship.
I have really been wanting to talk to you lately.
Our next date is going to be a long conversation.
We really need to talk about your behavior.
We need to talk about why you have not had time for me lately.
Can we talk about this.
I will give you a call.
We need to take some time out and really talk.
Maybe we should get together and talk okay.
We will talk.
Hey if we need to talk about something let us get together and discuss it.
Maybe we should talk about where we stand in this relationship.
If there is a change in feelings we can bring that out into the open.
I love for you to give me a call to chat.

Frame #24: Maybe I'm overreacting.

Maybe it is nothing.
I could be way off base here.
Maybe I am getting upset unnecessarily.
Excuse me if you think I am reading too much into things.
I hope I am not overreacting.

Frame #25: If you don't want to see me, just say so.

If something is wrong please do not beat around the bush and avoid me.
Tell me.
I would really understand more if you would just call me once and tell me what is going on.
Why do not you just tell me what you want to do.
If there is a problem just say so.
If you know you can not do something or do not want to just say so.
If you can not fit our friendship into your schedule I wish you would just say so.
If something is bothering you please tell me.
Just do not string anybody along.
If you do not have time for me just say so.

If you do not want to hang out with me let me know.
If you do not want to do anything anymore just tell me.
If you do not want to go out with me just tell me.
If something is bothering you tell me about it.
I would appreciate it if you just told me how you feel rather than avoiding me.
If you are just blowing me off let me know.
If I am not that important to you just tell me now.
Just say no.
If not tell me now.
If you do not want to do anything just say so.
If you do not want to go and there really is no meeting please just tell me.
If you do not want to see me I wish you would tell me instead of blowing me off.
What I really wish is that you be straight with me.
Tell me how you feel even if you think it is something that you think I do not want to hear.
If this is the case why did not you tell me.
If you can not just tell me now.
If you think otherwise then that is a decision that I will accept.
If you do not want to get together just say so.
If something is bothering you please tell me.
Terry if you just do not want to spend time with me then just tell me.
Will you please level with me.
Just tell me already for gods sake.
If you do not want to talk to me okay.
Hey if you do not want to hang out with me anymore that that is cool.
If you do not want to do things with me tell me.
I would rather have you do that than to have you always make plans that you cancel.
I will trust what you say.
Do you really want to spend time with me.
Hey just tell me if you would rather not do anything with me.
If you have a problem with me I would like you to talk to me about it.
If your life no longer accommodates our friendship then just tell me.
Do not just avoid me.
If so just let me know.
Do you want to spend time with me or not.
If you do not want to do anything just say so.
If not just say so.
If there is a problem please tell me.
If there is something bothering you I wish you would let me know.
If you do not want to hang around with me anymore just say so.

Frame #26: I want to have an honest talk.

I want to get this out in the open.
I want us to be completely honest with one another.

I will understand if you just talk to me and are honest about how you feel.
I would rather you be honest with me so I am not just wasting my time.
As a friend I hope you are able to tell me your reasons honestly.
I want you to be straightforward with me about this problem.

Frame #27: We can solve this problem.

Maybe we can work something out if you tell me what is wrong.
Maybe we can work it out before it goes any further.
We can work this problem out.
This problem can be resolved.
Let us try to work something out.

Frame #28: I'll find alternatives.

If you cancel I will just do something else.
I will find someone else to go with.
If something comes up for tomorrow night then I will plan on doing that.
I will call someone else who does not mind spending time with me.
Since you have a lot of activities I think it is time I get more involved and meet more people.
I will make plans with someone else.
I will go out with someone else tonight.
I will find someone else who has time for me.
I will find someone else to go.
I do have other friends I can go out with.
I am going to ask someone else to go.
If you have another unexpected meeting I will still have someone to go with.
I guess I will just call my friend and see if he wants to go.
If not I will try to find someone else with free time.
If you are going to have something else come up then I will probably go with someone else tonight.
I will find someone else to go with tonight.
Maybe I can find something else to do tonight.
I invited someone else to go with us.
I can make other plans for the evening.
I will just find someone else to go with.
I will go by myself.
I still have to find someone else to go.
Some friends called.
I am going to do something with them.
I will see if I can find someone else to go with tonight.
I might just call someone else and ask her to go out with me.
There will be someone.

I am going to get on the phone and see if I can find someone.
I have other friends too you know.
I am not sure if I will have something planned by then myself.

Frame #29: I'll stop bothering you.

I will leave you alone.
Well you do not have to worry about me anymore.
I will not bother you again.
You will not have to call and cancel anymore.
I will not bother you next time.
I will stop trying.
I give up.

Frame #30: You take me for granted.

I feel you have taken me for granted.
I will not be taken for granted.
Do not take me for granted.
You think that I will be by your side no matter what happens.
If you expect me to be always waiting around for you to do something you are wrong.
Now you take it for granted that you can blow me off.
I feel like you are taking me for granted.

Frame #31: I'm sorry for my behavior.

I really feel bad about putting you on the spot like this.
I do not want to make you choose.
I am sorry for the ultimatum.

Frame #32: Understand my perspective.

You are not being understanding to me.
Be more understanding about my feelings.
I hope you realize what you have been doing.
I do not know if you have experienced the same thing before.
Think of how you would feel.
You have to understand my situation.
Would you want to be treated the way you have been treating others.
How do you think it feels to be stood up.
That is the way I feel.

I mean you have to think a little of me too.
I would think that you would give me some type of consideration and concern.
I hope that you will give it some thought and put it into perspective.
I think you should try to understand how you would feel if you were in my shoes Terry.
I wish you would show a bit more consideration.
Just try to be fair to me.
Try to see my side of it.
Think about what I said.

Frame #33: If our friendship isn't important, then forget it.

I do not want to be your friend anymore if you are never going to treat me as if a friend.
If our friendship is not important then do not bother making sacrifices.
If you do not want to go out with me do not ask.
If those meetings are much more important than me then I do not want to be around you.
I think we should go our separate ways for a while til you get your priorities straightened out.
I do not feel like going out with you until you make a commitment to your friends instead of your little meetings.
If you can not make time for a close friend then just forget it.
Make some time to spend with me.
If your other commitments are more important than spending time getting to know me better then attend to those other commitments.
If I am not important do not lead me on and make and break plans.

Frame #34: If our friendship is important, then make time for me.

Make time for me.
As a friend you owe it to me to keep your commitments with me instead of throwing me aside every time something comes up.
Make an effort toward something we do together.
I think that you are able to make time for an old friend.
Work some time into your life for me.
I am not saying that every time I want to do something you have to drop everything and run right out.
Maybe you should start making some sacrifices if our friendship is important to you.
If you make a commitment to me I would like to think that I was important enough that you would not cancel out to go to some meeting.
All I am asking for is some of your time to enjoy a movie.
I think I deserve some of your time.
I feel you should make time for us to go out.
If our friendship means a lot you should not cancel our next date because of a meeting.
I feel I deserve respect.
Am I not important enough to make time for.

Make time for me.
Make an effort.
Find time or lose mine.
You should allow some time with them.
You have to arrange your time around your friends not just whenever you can fit them in.
If you want to maintain our friendship you would better arrange a time for friendship and a time for work.
If you want our friendship to last you will have to make sacrifices too.
If we are such good friends you should not keep blowing me off.
If this is the way you want to deal with it then I guess we will not be seeing each other again.
You should be able to set aside time for me.
If we want to keep it strong we need to spend time together.
That is if you actually care about me as a friend.
If our friendship means anything to you you will reconsider cancelling our date.
To do that you have to really want to spend time with me to.
I should be important enough to you that you do not cancel out on me all the time.
Please treat it as such.
If you did you would show a little more loyalty.
If you value our friendship at all you will be the one to make the effort from now on.
Why do not you just take a minute out of your busy schedule one of these days if you can.
Therefore if you want to continue being friends make time for us.

Frame #35: Friends are important.

Just remember in life friends are the most valuable people to you.
Remember friends are precious.
Friends are important too.
You have to realize that taking time out to have fun with friends is important too.
It is also important to take time out for your friends.
Friends can be if you try.

Frame #36: You should consider your priorities.

I suggest you start to think of others instead of just yourself a little bit more.
Obviously you need to evaluate our friendship.
You can not give one hundred percent dedication to everything under the sun.
I think you should think about how much our friendship means to you.
You have got to start thinking about your personal priorities.
Think about what is going to make you happy later in life being a super career success or having friends.
Think of all the commitments you are making.
Consider how much our friendship means to you.

Maybe you should spend a little time thinking about the way you have been treating other people lately.
You should get your priorities straight.
Is the meeting really more important than our friendship.
I think you need to decide your priorities.
If you can not get your priorities straight then I can not get you straight.
Maybe you should take some time out from your busy schedule to think about what is more important to you true friends or meetings.
Terry I think you need to get your priorities straight.
Think about how you have treated me your so called friend.
I do wish though that you would consider our friendship as important as all your clubs and meetings.
Maybe when you decide that I am a priority also in your life you should give me a call.
Sometimes you have to decide what is really important in your life.
Try to evaluate your priorities.
These groups and meetings are not going to be around forever.

Frame #37: You should make some compromises.

Now it is your turn to make concessions.
You have not met me half way.
You need to make more compromises.
A friend as any other relationship involves sacrifices.

Frame #38: I can't count on you.

I will not depend on you keeping plans with me any more.
If you always cancel engagements how can you expect to be there when I really need you.
I do not even count on you anymore.
I do not trust you to be able to free your schedule.
All trust and faith I had in you is gone.
Terry I do not believe you.

Frame #39: You need to make choices.

Pick some that are important.
Leave the others alone.
It is your decision.
So what is it going to be boys or not.

Frame #40: You have made me a low priority.

Now it seems that your work and other commitments are more important than our friendship.
Lately it seems like I am the last thing on your list.
Lately Terry you have placed our friendship second to your meetings and activities.
Obviously it is not a priority of yours.
It seems like I am always last on your list.
It is becoming quite obvious that these activities and everything else come before me.
Maybe I do not come last.
It is hard to always be second to your meetings.
You do not give me anything.
You put everyone else before me.
It appears you never have any time for me but plenty of time for others.
I can understand having a few priorities that come before me but not everything.
I am not as important to you as you are to me.
Maybe I do not mean enough to you to keep a simple bowling date.
Am I not as important as all these other things.
It seems to me that we get together whenever it is convenient to you.
Something always comes up that is more important that me to you.
I am feeling unimportant in your life.
It seems to me that these activities of yours are more important than our friendship.
You have been seeing another guy.
Whenever I try to plan something for each other you have something more important to do.
It seems to me that our relationship is not very much of importance to you if you can only back out of ourengagements.
Obviously this relationship is not that important to you.
These meetings and things must be more important than our friendship.
You can not break yourself away from them ever.

Frame #41: I make time for you .

I will not blow you off as much as you have blown me off.
Somehow I find the time to make room for the important things in my life like friends.
I make time for you.
I manage to make time for you without canceling at the last minute.
I set aside time to be with you tonight.
I have adjusted my schedule to keep the time open.
I have been giving of my time to you making an effort to see you.
I can to find time for you.
I was able to set aside time for you.
You know I always make time for you.
I know that if I got busy I would always manage to find some time for you.
I make time to spend with people that are important to me.

Frame #42: I have a busy life, too.

I am busy too.
So am I.
I am busy too.
You know I have a busy schedule too.
My life does not revolve around you.
I have other things to do.
I am also involved in other things.
I have other friends and activities that I may have forgone because of her.
I have other things to do.
I have social commitments also.
I do not have the time to play games.
I will be pretty much busy whenever you call.
Try to realize that I have a life to.
I am busy too.
I really need to get going.
I need to make plans too.
My time is as valuable as yours.
I have things that need to be done and activities to do also.
I have a busy schedule also.
I have other things to do you know.
I am busy.
I have too.
You are not the only one with a life you know.
So do I.
We are both busy people.
So do I.
So do I.

Frame #43: I've made sacrifices to be with you.

Every time we make plans I spend time getting ready.
When I set aside time to go out I am canceling potential activities.
I keep canceling plans to go out with you.
I have broken plans with other people to be with you.
I have made concessions before for you and the sake of our friendship.
Sometimes I miss meetings or other things so we can go out.
I broke other plans for tonight.
I have made sacrifices in an attempt to do things with you.
I have canceled appointments in order to go out tonight.
I always miss out on other opportunities.
When you cancel plans that sometimes means I miss out on other things I could have dome.
I have gotten a lot of work done so I could go out tonight.

I am taking time out for you to do these things.
Then I have to keep changing my schedule around as well.
That way I will not plans around something when I could be getting work done.

Frame #44: I'm not angry.

Those are not the feelings I have been having.
I will not be upset.
I have been getting this strange feeling lately.
I am not mad.
I am not mad at you.
Do not take this the wrong way.
I will not feel bad either.
Do not think I am being too harsh.
I am not mad.
I will not be mad.
I am not disappointed.

Frame #45: I have tried to set up dates with you.

I have been patient with you.
I have tried my best with you.
That is why I keep planning things for us to do.
I have been trying to plan things for us to do for a while now.
I have been trying to plan things for us to do for a while now.
We will try to organize some kind of activity.
I try to be a man of my word and keep promises.
I do not know why I even bother making plans with you.
I am trying so hard to be friends with you.
Terry I have tried to call you and make numerous appointments of late.
I have tried and tried to keep our friendship up.
I have made a constant effort to se you.
I have tried to be reasonable and understanding.
It seems that I am doing all of the work to try to keep it that way.
I am trying to be understanding.

Frame #46: Since you cancel, I won't make plans with you.

We will stop planning things.
I do not plan to go and decline other activities to be stood up.
I am not going to make tentative plans only for you to cancel them.
I do not want to keep making plans if they are just going to fall through.

I do not want to make plans and have you cancel again.
I will not be making plans with you.
I just can not keep rearranging my plans all the time like this.
Terry I do not want to make a date with you anymore.
I do not want to make plans with you anymore.
I would rather not try to make plans if you are going to keep canceling them.
I did not really want to go out anyway.
Mom always told me about scheduling my life around incompetent irresponsible people like yourself.
I do not want to get all excited about going tomorrow night only to have you cancel.
Since you are not able to find time for us neither can I.
I do not believe in scheduling my time to do activities with her when you cancel at the last minute.
Since I am continually put in this situation I am not going to plan anything with you anymore.
I can not revolve my life around your schedule.
I do not want to have to plan around your schedule.
I am not willing to exclude other events when all you will do is cancel.
I do not want to say yes to tomorrow if there is a chance you will cancel.
You will not hear from me in the meantime.
I am not going to wait around for you to eventually follow through on your plans.
I do not want to get my hopes up just so you can tear them down again.
I can not afford to keep making plans with you and then having you cancel out all the time.
If I back out on you at the last minute like you have done the last ten times do not be surprised.
You are only getting a taste of your own medicine.
I do not want to schedule my life around your meetings.
I am not going to go out of my way to make up stuff to do.
Until then I would rather not keep reserving time for you to have to cancel at the last minute.
If you keep blowing me off I may choose not to do anything with you anymore.
I am not going to stand for this anymore.
I will not be able to make my plans around you.
I just do not have time to sit around and wait for you to cancel plans.
You know if you are just too busy to get together then let us just not bother making any plans for a while.
The more you cancel the less I want to go out.
It is too hard to make plans with you.

Frame #47: If you don’t want dates don’t make dates.

I wish you would not make plans to try and make our relationship better and then make things worse by breaking our date.
I mean why make plans when you cancel all of them.
If you have a problem going out with me then stop making plans with me.

If you do not want to get together tomorrow do not bother.
If you really do not want to go out do not make excuses.
You should not make plans you can not keep.
If you do not want to do anything then do not make plans in the first place.
If you do not want to do something we do not have to.
Now if you do not want to date me then just stop calling and making other arrangements to see me.
Terry if you do not want to hang out with me then do not.
Do not keep making plans.
Do not make any plans with me that you can not keep okay.
Do not make plans when you are more than likely to back out.
If you do not have time do not ask me.
If you do not have the time then do not make the plans with me.

Frame #48: I expect you to cancel tomorrow, too.

How do I know you will be there tomorrow.
I do not know that if we make a date you will keep it.
Can you go to the movie tomorrow night.
Tell me now if you are going to go.
I am sure as far as you are concerned it is not set in stone.
Are you sure you do not have something else planned for tomorrow night.
Are you going to puss out on me again.
Can it be my turn to do what you do tomorrow night.
Will you be there.
Are you sure you can make it.
Are you going to cancel out on me again.
You will probably be doing something else.
What happens tomorrow when another meeting just pops up.
Why do not you check your calendar to see if tomorrow is okay.
Do not you have any meetings tomorrow.
The question is will you go.
Jesus you sure nothing else will come up tomorrow.
Are you sure you do not have anything else to do.
You would probably just cancel that date too.
You can cancel them and get it over with.
If you can still go tell me now.
Terry are you sure you can make it tomorrow.
Sure you will not cancel on me tomorrow will you.
Making a date for tomorrow night might be too difficult a commitment for you.
Are you sure you will not have anything else to do tomorrow.
Terry are you going to call me in an hour and have an excuse that you can not go tomorrow night.
I am so accustomed to you canceling.

You will cancel.
Something may come up once again.
I kind of expected this to happen.
Are you going to cancel again.

Frame #49: Call me when you have more time.

If you do not have the time to spend with me to go to see a movie then call me when you are sure you can go out.
Just call me when you find yourself with nothing to do.
When or if you figure that out I will think about doing something with you again.
Why do not you call me when you can fit me into your precious schedule.
Why do not you give me a call sometime when you are free.
Call me back when you can definitely go out.
Tell me if you want to go out.
Call me when you want to go out.
Why do not we just get together when you have more time.
Give me a call when you really have time and are ready to keep your commitments.
You call me if you want to do something with me.
Give me a call when you have the time.
Give me a call when your busy schedule has an opening and you decide I am important enough to fill it.
Give me a call when you want to get together.
Give me a call when you have some free time.
Call me when you think you can be a responsible friend.
Call me when you get a chance.
Why do not you call when you have some free time for sure.
Maybe someday you will discover that you have a few extra minutes to spend with your friend.
From now on you call me when you have some free time.
Why do not you call me when you will not have anything else to conflict with our time all right.
When you get home from this important meeting why do not you call.
Whenever you can really do something call me right before.
Next time it works for you call me.
Let me know when you can for sure do something.
Well when you do find time and want to get together give me a call.
If you have some time one day give me a call.
If you want to go out tomorrow call me.
Call me when you have some free time.
Pencil me in for june.
Call me at your convenience.
From now on Terry you make the plans.
It is up to you from now on if you want to keep this friendship up.

Call me when you are free sometime.
When you decide you want to see me pencil me in on your agenda.
Call me.
When you have time to do something for sure then we can talk about it.
Call me when you have some time to see a movie.
If you want to do something that you will not cancel out on give me a call.
When everything calms down for you call me.
We will go out when everything calms down for you.
Give me a call.
This time we will work around my schedule.
When you decide you have some time call me.
Maybe your schedule will not be so booked up in the future.
Call me sometime when you can fit me in your schedule.
If you find yourself in a position which someone stands you up give me a call.
You have my number.
Let me know when you have time for me.
Call me when you truly want to get together.
When you get a chance look at your calendar.
Give me a call.
Listen why do not you call when you can guarantee you have some time to actually do something with me.
Maybe next weekend we can do something if you have time.
So call me when you do.
You will just have to call me sometime when you are free.
Just call when you decide you would like to do something.
Why do not you call me when you are sure you do not have anything else.
Call me when you actually want to get together and not just talk about getting together.
Call me with your answer if you can find the time.
From now on call me when you really want to do something.
Then when your schedule allows maybe we can do something.
Why do not you just get back to me when you are sure that you will have time to do something.
Perhaps when your schedule lightens up we can get together on the spur of the moment to do some things the way we used to.
So when you get a chance to fit me into your schedule let me know.
Call me in a few days when you are sure you have the time.
Terry why do not you give me a call when you are absolutely sure you can follow through.
If you think you can make time for me someday give me a call.
Just give me a call next time you are sue you have the time to do something.
As a matter of fact do not bother to make any plans with me in the future until your life is less complicated.
When you think you have real sufficient time well than maybe give me a ring.
From now on it is going to be your obligation to make plans with me if you feel you want to.
When you are sure you have time to fit me in then you call me.
Plans can be carried out much easier if you become the initiator.
Just give me a call some time when you can actually make plans and stick with them.

When things slow down we will get together.
Call me when you have a chance.
If you ever find time to actually do something with me give me a call.
Just give me a call if you want to.
When you really know or actually have time to be with me or do something with me then give me a call.
Call me when you can fit me into your schedule.
Please call me whenever you have time.
Why do not you call me for plans when you can fit me into your schedule.
Tell you what when you have time you give me a call.

Frame #50: Let's talk tomorrow about the date.

Call me early to confirm our date.
Why do not you give me a call tomorrow night.
I will call you to confirm our date.
If you actually can make the movie give me a call tomorrow.
Why do not you call me tomorrow and let me know how your agenda looks.
Call me tomorrow.
Give me a call tomorrow night.
Call me later to confirm tomorrows date.
Give me a call tomorrow night to make sure I am free.
Call me tomorrow.
Just call me tomorrow.
Why do not you call me tomorrow around six if you still want to go.
Why do not you call me tomorrow and find out.
Call me tomorrow to let me know if you can go to the movie.
I will have to call you tomorrow to let you know about the movie.
I will give you a call tomorrow and see what is up.
Why do not you call me tomorrow when you are sure nothing else is going to interfere.
About the movie tomorrow night call me tomorrow night.
Maybe you should just give me a call tomorrow.

Frame #51: Keep our date tomorrow.

No excuses.
This time stick to your word.
I do hope you will be sincere about your commitments.
I just hope that you consider me if other engagements arise.
We better go to the movie tomorrow.
Just do not blow me off.
Do not cancel out at the last minute.
If you go out and do something with me you make the arrangement this time and stick to it.

Please make sure to keep our date for tomorrow.
If you really have all these meetings be more responsible to me.
Are you sure you can go tomorrow night.
Do not be breaking plans with me.
Terry I wish we could set up a time to do something and have it stay.
Just do not tell me you can go and then cancel again.
Please try to show up tomorrow.
Keep future commitments to me please.
Do not call me up tomorrow and tell me you forgot and planned something else okay.
I do not think it is too much to ask that you occasionally follow through with a commitment.
Just do not cancel on me okay.
It would be nice if you would keep a few of our engagements.
You better show up.
Are you sure you can go tomorrow night.
Try not to keep cancelling things on me.
Hey next time we make a plan let us try to stick to it.
If you really want to do something find a time and stick to it.
Next time do not make plans and cancel out.
You have to learn to stick to prior commitments.

Frame #52: I hope we get together tomorrow.

I do hope you will be at the movie tomorrow.
Hope to see you soon.
I hope you are able to make it.
I hope you will join me.
I hope you can find time in your busy schedule to come with me.
I hope things work out.
Well hopefully I will see you tomorrow.
Hopefully I will see you tomorrow.

Frame #53: I understand your constraints.

I understand you miss dates.
I understand that you have other things to do.
I have tried to understand.
I understand you have obligations and responsibilities.
Okay Terry I understand.
If you have things to do you have things to do.
I understand that you are busy or whatever.
I understand if you are too busy to go out.
I guess meetings are pretty important.
I have tried to be understanding about your hectic schedule.

I understand you have a lot to do.
I understand if you have a meeting.
I realize you have to go places and do things.
I understand that you have other commitments.
I understand that things do not always work out.
I could understand if you were out with a gal.
Terry I realize you have a lot of other obligations.
I realize we are both very busy.
I understand.
I will understand.
I am not saying you cancel for the wrong reasons.
I realize you have a busy schedule.
I have been trying to understand that you have been busy lately.
If you honestly think you will not be able to go tomorrow then you do not have to take me.
I realize that your activities are very important to you.
I understand that you have a meeting to go to today.
I accept that you have a meeting tonight.
I know you have a meeting tonight.
I understand your other obligations.
I realize this meeting is important.
Your meetings and group activities are important to you.
If you have to go to meetings fine.
I understand that you have meetings and everything.
I can understand that you have other things in your life.
If you really do have a meeting then I guess I understand.
I know you have lots to do.
I understand that there are a lot of worthless meetings we have to attend at school.
I understand this meeting arose suddenly.
I understand about the meeting.
I understand that you have meetings.
I understand.
I realize you have other obligations.
I understand that you wish to succeed and do well in the world.
I understand you have commitments.
Look Terry I understand that you have a lot of things to do and a lot of commitments to keep.
I realize you have a lot to do.
If the only reason behind your last minute changes is that you have gotten involved in a lot of organizations that take up much of your free time I understand your motives.
It is not the things you have to do meetings commitments and such.
I know you could not help it this time.
I am not saying it is always your fault either.
I will try and be understanding toward your commitments.
Terry that is too bad about your appointment.
I know people.
Their tastes change.

Well Terry I understand your situation.
I understand that your meetings are important to you.
I can understand breaking a commitment once in a while.
I realize that good friends can get away with something here and there.
If you can not meet me you just can not meet me.
I understand that you have a lot to do.

Frame #54: Sorry, I can't go tomorrow.

No Terry I will not go to the movie with you tomorrow night.
I should tell you I can not see the movie tomorrow.
No I do not think I can make it tomorrow night.
I am busy tomorrow Terry.
I can not see the movie.
No Terry.
No Terry we can not make it tomorrow night.
Sorry I do not think we should plan on going tomorrow night.
I am sorry.
I can not go to the movie with you.
I have got plans tomorrow night.
No I can not go tomorrow.
Sorry Terry I can not make it tomorrow night.
I am sorry.
I already have plans.
No Terry I have already made other plans.
I doubt I can make it tomorrow.
No we can not go to the movies tomorrow night.
You know Terry I just remembered I have a meeting tomorrow night.
No I do not think so.
No Terry I will not go to the movie tomorrow.
I may be busy tomorrow.
I do not want to go to the movies tomorrow night.
I promised someone else that I would do something else tomorrow.
I think I have a meeting tomorrow night.
I think I have a meeting from now on when you want to do something with me.
No Terry tomorrow night is not good for me.
I am sorry.
I do not have time for broken dates.
No Terry I am sorry.
No.
I have plans tomorrow night and every night for a while.
I think if we went tomorrow I would be in a cranky mood.
I am sorry Terry tomorrow night will not work.
No I am sorry.

Tomorrow night I need to study for that exam I told you I have.
I might have an emergency underwater basket-weaving meeting tomorrow night.
No I already have other plans.
I am sorry Terry.
I have plans tomorrow.
I will not be able to make it.
I am sorry I can not.
I am busy tomorrow.
No I will be busy tomorrow night.
I am busy tomorrow night.
I can not make it tomorrow.
I am busy tomorrow night.
Terry I do not think I can go tomorrow night.
Maybe we better not go to the movie tomorrow night.
I am sorry.
I am sorry tomorrow night is not good for me.
I am already busy tomorrow night.
No I can not go tomorrow.
I have plans for tomorrow already.
It is too late now.
I can not make it tomorrow.
I am not sure if I will be able to make it tomorrow night.
I do not have time to go to a movie tomorrow.
Sorry Terry I am busy tomorrow.
Terry I am sorry.
Sorry Terry I am going to the movie with jane.
I am sorry Terry.
I just can not.
No that is okay.
Tomorrow night no I do not think so.
However I can not make it tomorrow.
I have already made plans for tomorrow night.
No Terry we can not go to the movie tomorrow night.
I have an important meeting I have to attend.
No we can not go tomorrow night.
I am busy tomorrow.
Well Terry I am kind of busy tomorrow night.
No I do not want to go tomorrow night.
I can not go tomorrow.
I have to do some other things.
Hell no we can not go tomorrow.
I have plans tomorrow night.
I am sorry Terry.
I have a previous engagement on for tomorrow.
I was just going to call and tell you I have to cancel tonight's plans.
I will be busy tomorrow night too.

Sorry can not make it.
I do not think I can go tomorrow night.
No I do not think that would be a very good idea.
I am not going to say yes we can go to the movies tomorrow.
I am not going to go to the movie with you tomorrow.
Tomorrow I do not think so.
You see I am going out with some real friends who stick to their commitments and do not cancel out on me all the time.
No I have plans for tomorrow night.
Sorry kid.
No Terry.
No I think I will be at meetings all night and day.
About going tomorrow I can not.
No Terry we can not go to the movie.
No I am busy tomorrow.
No Terry this is bullshit.
I have already got plans for tomorrow night.
I can not make it tomorrow.
I do not know if we can Terry.
I am not sure what is going on tomorrow night.
I think I am going to go to the movie with a date.
Terry I am sorry.
I am sorry you are going to miss the movie tonight.
Well I am busy tomorrow Terry.

Frame #55: If you cancel again, I'll be angry.

If you cancel out tomorrow so help me I will kill you.
I will not make an effort to go out if you miss tomorrows date.
I am really going to be pissed off if you cancel again.
I will be honked off if you cancel again.
If something comes up for you tomorrow I will be more than upset.
If you are not there tomorrow night I will never believe you again.
I will give you one more chance.
You will have some explaining to do if you do not show up.
We will not be good friends anymore if you do not show up.
If you do not show up then I can not be your friend anymore because it is too painful being let down.
Just do not commit to do something with me and back out again.
I really miss talking to you and would be very upset if you have to cancel again.
In the future however I do not want to see this happen again.
This time I will take no excuses.
I do not want to have to cancel again.

Frame #56: Okay, I can go tomorrow.

It is okay.
Well Terry I can go to the movie tomorrow night.
Sure I guess tomorrow is okay.
Tomorrow night I am going to the movie.
Do not worry about it.
That is okay.
Well we can make plans for tomorrow.
Yeah I guess.
Fine we will see the movie tomorrow.
Yeah.
Well I guess it is okay.
Okay.
I have no plans for tomorrow night.
Okay we can go tomorrow.
I am free.
Yeah sure.
It is all right if we change it to tomorrow.
Okay sure.
Terry it is okay.
Sure no big deal.
Do not worry about it.
Okay.
We will try to go tomorrow.
Okay.
Sure I will go.
Sure.
See you tomorrow.
Hey whatever.
Do not worry.
Ok tomorrow sounds good.
I will go tomorrow.
Okay.
I would go to see the movie with you.
Fine.
Yes we can go tomorrow night.
Yes.
Fine Terry I do not think I have anything on for tomorrow night.
Yes let us go to the movie tomorrow.
Tomorrow night is fine.
That is ok.
I guess we can go tomorrow night.

That is ok.
Okay.
That is fine Terry.
We can postpone til tomorrow night.
All right I could wait until tomorrow to see you.
Yes we can go tomorrow night.
That is all right.
Well I suppose.
Sure I will go tomorrow to the movie.
I guess we can go to the movie tomorrow.
That is fine.
It is all right.
That is okay Terry.
All right.
Yeah sure tomorrow is fine.
We will see each other tomorrow night.
Fine.
That is fine with me.
Tomorrow is all right.
Okay whatever.
Yeah I suppose we can.
I would love to move it to tomorrow night.
It is no big deal.
It is all right.
Sure we can go tomorrow night or the night after that or the night after that or the night after that.
We can go whenever you want.
Better yet I will just pencil you in for my calendar for tomorrow night.
I guess if you can not go then I will try to see you tomorrow night.
Yes we can go to the movie tomorrow night.
Terry that is all right.
I guess I will go with you tomorrow.
That is okay.
I guess we can go tomorrow night.
Okay we will go to the movie tomorrow night.
Okay.
Of course I can make it tomorrow night.
See you tomorrow.
Well let us try for tomorrow.
Sure.
Yeah I guess.
No problem buddy.
I am willing to go to the movies tomorrow night.
Yeah tomorrow night is fine with me.
That is okay.

Frame #57: Tomorrow is fine if we can make it.

If we are both free we will see the movie.
If you are going to go with me that is great.
If there is even a small chance that you think you can not tell me.
If you really want to and you are not busy we can go tomorrow.
If you really have been busy and want to go out tomorrow fine.
If you are available tomorrow I would love to go.
So long as you are not lying to me or anything I think I am all right.
I will see you tomorrow unless I am busy that is.
Gosh Terry I do not know if we can or not.
If you really want to go out with me tomorrow fine.
That is fine if you can make it tomorrow for sure.
If you want to get together tomorrow that is great.
If I have nothing else to do tomorrow I will go with you Terry.
If you really want to go to the movie tomorrow night we can.
Maybe we will do something.
Terry I will give you till tomorrow night.
If you do not think you can go tomorrow we will not.
Why do not we go ahead and make some plans now.
If I am free I will go with you.
If you are sure you can go tomorrow fine.
If you have your meeting and it is important I guess it is okay if we go tomorrow night as long as you do not have any reason to cancel that date.
We will see what happens.
I will go to the movie tomorrow night if you know absolutely that you will be there.
I will plan another meeting if you promise to be there.
If you can find the time real time to meet with me even if just to have a coffee I would be delighted.
Well do you think you can make it tomorrow.
I will give you a chance tomorrow night only if you are sure no other meetings will come up.
I can not count on the fact I can go.
I will go if I have not made any plans.
If you want to tomorrow fine.
As far as tomorrow I will only make the plans if you are sincere about keeping them.
I want to be sure that you really want to go.
That is fine if you definitely can do it tomorrow.
Maybe we will do it again tomorrow night.
If I am home I will go with you.
Well Terry do you really plan on going tomorrow.

Frame #58: Okay, go ahead and cancel.

It is fine if you can not go to the movie tonight.
Terry go ahead and cancel.
Go ahead with your meetings.
Sure blow me off one more time.
Otherwise it is to be like this all the time okay.
Oh really that is just great.
Sure blow me off again.
Go to your meeting.
Since you have been so busy I think it would be more important for you to keep your engagement.
It is really not important.
You just go to your meeting.
Go ahead.
Go to the meeting.
Okay Terry whatever.
Okay if you really have to go.
If the meeting seems more important this time that is okay.
Thanks for calling and telling me you could not go instead of just not showing.
Fine just go.
If it is sheer coincidence then okay.
Do not worry about it tonight.

Frame #59: Check with me another time.

We will see if I can go.
Maybe some other time.
Call me another time.
I will see what I can do.
I will check my schedule.
I can not promise I will be free.
I will see if I can go.
Maybe we can go sometime next week.
I will adjust my schedule.
Maybe I will be able to see you in the future.
Let me check my schedule.
I will get back to you.
If you ever for any reason should want to call me you can.
We will try to plan something for next week.
Call me later.
Maybe some other time.
Just give me a call whenever.
You can call me next week.
Maybe I will have time to go.
I will try to fit you in my schedule.

From now on you had better call me to do something and then check to see if I am busy.
Maybe we can do something another time.
Why do not you give me a call next week or something.
Maybe some other time.
If you want to do something with me you call me and make the plans.
I will see if I can fit you in.
I will let you know if I can fit you into mine.
Maybe we can set a date to do something.
If I do not have any plans I might consider it.
The next time we make plans it will be when I am sure I have nothing better to do.
If you ever want to have any contact with me in any way you will have to call or come over and see me in person.
My schedule is far more flexible.
If you want to do something give me a call sometime.
Maybe some other time.
Why do not you call me and decide.
I will try to fit you in.
We can plan an event or date then.
Perhaps we will get together soon.
Even better maybe we can work out a good time to get together.
I will see if I can pencil you in.

Frame #60: You leave me with nothing to do.

This way one of us will not be waiting on the other if something comes up.
Why do you always leave me hanging like this.
Why do you put me in the situation of always waiting around for you to call me.
I have wasted my time waiting for you.
Your constant canceling infringes on my time.
It is not fair for me to sit home with no plans.
I should not have to risk my night wondering if you will cancel.
I have had so much free time lately.
I am almost always free.
You will cancel again tomorrow and leave me with nothing to do.
When you cancel I think of what I could be doing other than sitting at home.
That leaves me with nothing to do.
I do not want to wait around.
Time is money.
You are causing me to waste a lot of time.
You always leave me with nothing to do.
I do not need to sit around and worry about your canceling.
Now I am left alone.
I am stuck with empty plans.
I have plenty of free time lately anyway.

You keep changing our plans and leaving me with nothing.
You have left me without anything to do too many times.
Because of you I am unable to do any of these things.
You are wasting my time.

Frame #61: If you called sooner I could have replaced you.

At least you could call me.
Try to tell me your plans a little sooner.
It would have helped if you had called sooner.
I could have replaced you.
There are other things I could have planned if I knew you were going to cop out on me.
I wish you would find out about these meetings sooner.
It is difficult for me to plan other activities on such short notice.
I can not just be told on a whim that things are canceled.
Well it is a good thing you called.
You know if you had just told me three days ago you could not go tonight I would not be upset.
It is the way you wait until the last minute to tell me.
I could have made other plans.
I could have been going out with other friends.
It might be kind of hard to find a date at this time of the day.
If you have to cancel you should not cancel at the last minute.

Frame #62: We had a date.

I have plans.
Do you know what a previous engagement is.
Terry we made these plans to attend this movie several days ago.
We had plans.
We had a meeting of our own planned for tonight.
Look Terry we already had plans for tonight.
I have plans.
We planned on tonight.

Frame #63: Skip the meeting and keep our date.

Put everything else on hold.
All you have to do is cancel the other engagement to go out with me.
Instead of canceling our plans call and find out what you are going to miss.
I am sure you do not have to go to every meeting.
Could you please consider saying I can not be there I have a previous engagement and get the

information the next day.
I think you should cancel the other meeting.
Break the other commitment.
Can not you just miss your meeting for once.
If you can try to find a substitute for the meetings.

Frame #64: What is the story on the meeting.

Do not you know about these meetings days before.
What meeting do you have to go to.
What meeting.
What meeting.
What meeting is it.
Terry could not you have scheduled your meeting at a different time.
What do you mean you just found out.
If the meeting is that important you would have known about it before.
Why did not you tell me about this meeting when we made plans for the movie.
Why is the meeting so important.
Why did you only find out about the meeting today.
What is the meeting for.
How long have you known about the meeting.
How do all these things suddenly pop up out of nowhere.
What is the meeting for.
Jeez Terry how many meetings do you go to normally.
What meeting do you have.
If you new in before that you had this meeting tonight why did you agree to go to the movies.
When did you find out about this meeting.
Please do you have to go to the meeting.
Why is it you can not arrange these things earlier.
What is the meeting for.
Do you really have to be there.
What kind of meeting is it anyway.

Frame #65: I want to see the movie tonight.

I was really looking forward to going to the movie with you tonight.
I want to see the movie tonight.
I was looking forward to the movie tonight.
I would really like to go to the movie.
I was really looking forward to tonight.
Since I have set aside this time I would like to take it.
I am going to the movie tonight.
If you want to see me you can see me tonight.

I am still going to the movie tonight.
I think we should go to the movie tonight.
I think we will go ahead and go.
I have been looking forward to seeing this movie for a long time.
I am still going to go tonight.
I want to see this movie tonight.
I had it planned for tonight.
I wanted to go tonight.
I am going to see this movie with or without you.
Well I was really planning on going tonight.
I am going to the movie by myself.
I am going to the movie anyway.
I was really looking forward to this movie tonight.
Terry I am sorry we can not go to the movies tonight.
I am really in the mood to see a movie tonight.
I really want to go.
I was looking forward to going to the movie.

Frame #66: When we make dates make sure you are free.

Do not make plans you will break.
At least keep commitments you make some of the time.
I think if we want to make any plans in the future you better make sure you are free on that date.
Make sure you have no previous engagements.
You should only make commitments you know you can keep.
Next time we plan to get together make sure you can make it.
From now on do not set dates with me when you have other commitments.
If you want to go out with me check to see when all your meetings are and work around them.
If you can not plan something and carry it through then do not plan it.
Before you make any more plans with me why do not you check your schedule over again.
From now on if we do anything it is going to be because you were persistent in asking me.
You pick a time convenient for you.
Do not bother calling until you are sure you can make the date.
If you do not have time to go out do not make plans.
That way you will not have to feel bad when you cancel out again because you will not be canceling out.
How about we get together when you are positive you have nothing else to do.
Are you sure you are free tomorrow.
Well I just wanted to make sure.
When you are sure you do not have other things to do I will make plans with you.
Do not make plans with me anymore unless you can keep our promises and dates definite.
Please before we make a date for tomorrow be sure you can make it.
From now on if you make a commitment to me you either keep it or else we just will not go

out any more.
You should know that when I plan to do something I would like to follow through to the end.
If you want to spend time together then plan ahead and make sure you will be able to do so.
Next time make sure of all of your prior engagements before you plot one for me.
I do not want you to make plans with me unless you are positive that no other commitments will interfere.
Please make sure you do not have any other meetings you have to go to before telling me we will go out and do something.
If you are not sure about a certain activity or something let me know at the outset instead of letting me get excited only to have the wind knocked out of me.
Next time you want to make plans do not make them unless you can keep them.
Do not make plans with me if you are unsure about being able to keep them.
So let us set a time when we can both definitely do something.
If we are both free sometime maybe we can do something then.
Let us make sure we make plans when we are positive there will not be any conflicts.
The next time you want to make plans with me you should really make sure you have nothing else to do.
If you want to do something with me the check your schedule and make sure you are free.
That way we will both be able to relax and ourselves.
Now are you sure you can make it tomorrow.

Frame #67: Be more organized.

Did you lose your I-book.
Try to find out about your meetings ahead of time.
Please next time write it on the calendar.
Do not you have a set calendar of events.
You need to get organized.
If you had an appointment book you would have realized that you had a meeting.
I just want you to become more organized.
Maybe you should also learn to budget your time and have some free time to explore friendships rather than constantly finding a meeting that may or may not exist.

Frame #68: Let's find a compromise plan.

Since your meeting is not til later let us go out for a walk right now and talk.
We can go to the late show.
Let us decide and go out at a time we both can make it most definitely.
Let us make a deal.
Come over after your meeting no matter what time it is.
If it is not too late we will go to a movie.
If it is too late we will just visit.
Maybe instead of going to a movie we could go to dinner.

I will see you for dinner tomorrow.
I will make you a deal.
Could I go to the meeting with you tonight.
I think it would be fun to attend your meeting.
I will be at the pizza place after the movie if your meeting let's out early enough.
You can meet me after your meeting.
We should go to the later show or the next day.

Frame #69: Goodbye.

Talk to you later.
Talk to you later.
Goodbye.
Thanks for calling.
Bye.
Have fun at your meeting buddy.
I appreciate your calling.
Let us not get into details over the phone.
Bye.
Bye.
Have fun at your meeting.
I have something I have to get done before I go out tonight.
I will talk to you later.
See you later.
I will give you a call later okay.
Well I have to go.
Bye.
Cool okay later.
Goodbye.
Have a good time at your meeting.
Bye.
Terry I will talk to you later.
Later.
Have fun at the meeting.
Have a nice day.
Talk to you later.
Thanks for calling me.
Have fun at you meeting.
See you.
Have fun at your meeting.

Frame #70: I don't know what to think.

I do not know what to say.
I feel in the dark here.
I can not read your mind.
I am not sure.

Frame #71: Do not feel obliged to me.

Please do not feel an obligation to go out of your way for me.
I think it would be better if you took time out from something when you wanted to rather than me pulling you away from it.
You do not have to call me tomorrow night if you are busy.
Do not make new plans just because you felt bad about breaking previous ones.
I know if you broke down and decided to see the movie you would not have a good time and feel forced into it.

Frame #72: Psychological prefaces.

Can you spare a few minutes.
You think about that.
Well why.
I do not know.
Terry can I tell you something.
I just want to tell you.
I just thought it was necessary to mention it and to tell you what was on my mind.
I hope you understand.
How does that sound.
Terry do you have a minute.
Do one thing for me.
I will tell you one thing.
Do me a favor.
That is all I ask.

VITA

Bruce Laurence Lambert was born in Gainesville, Florida on April 9, 1965. He graduated from Deerfield High School in Deerfield, Illinois in May of 1983. In January of 1987, he received his bachelor's degree in Speech Communication with honors and high distinction from the University of Illinois at Urbana-Champaign, and his master's degree in Speech Communication from the same institution in October of 1988. He has served as a teaching and research assistant at the University of Illinois at Urbana-Champaign and as a Lab Associate of the Knowledge-Based Engineering Systems Research Lab from 1987 to 1991. He is a member of Phi Beta Kappa. He currently teaches interpersonal communication and health communication as an Assistant Professor of Pharmacy Administration at the University of Illinois at Chicago.