

THE PSYCHOLOGY OF LEARNING AND MOTIVATION

Advances in Research and Theory

EDITED BY GORDON H. BOWER

STANFORD UNIVERSITY, STANFORD, CALIFORNIA

Volume 21



ACADEMIC PRESS, INC.

Harcourt Brace Jovanovich, Publishers

San Diego New York Berkeley Boston
London Sydney Tokyo Toronto

COPYRIGHT © 1987 BY ACADEMIC PRESS, INC.

ALL RIGHTS RESERVED.

NO PART OF THIS PUBLICATION MAY BE REPRODUCED OR
TRANSMITTED IN ANY FORM OR BY ANY MEANS, ELECTRONIC
OR MECHANICAL, INCLUDING PHOTOCOPY, RECORDING, OR
ANY INFORMATION STORAGE AND RETRIEVAL SYSTEM, WITHOUT
PERMISSION IN WRITING FROM THE PUBLISHER.

ACADEMIC PRESS, INC.

1250 Sixth Avenue

San Diego, California 92101

United Kingdom Edition published by

ACADEMIC PRESS INC. (LONDON) LTD.

24-28 Oval Road, London NW1 7DX

LIBRARY OF CONGRESS CATALOG CARD NUMBER: 66-30104

ISBN 0-12-543321-2 (alk. paper)

PRINTED IN THE UNITED STATES OF AMERICA

87 88 89 90 9 8 7 6 5 4 3 2 1

CONTENTS

AN INTEGRATED COMPUTATIONAL MODEL OF STIMULUS-RESPONSE COMPATIBILITY AND PRACTICE

Paul S. Rosenbloom and Allen Newell

I. Performance: Compatibility and Goal Hierarchies	3
II. Learning: Practice and Chunking	26
III. Discussion	43
References	49

A CONNECTIONIST/CONTROL ARCHITECTURE FOR WORKING MEMORY

Walter Schneider and Mark Detweiler

I. Introduction	54
II. Traditional Views of Short-Term Memory	56
III. A Connectionist/Control Architecture for Working Memory	57
IV. Interpretation of the Working-Memory Literature	71
V. Context Effects, Proactive Interference, and Release from Proactive Interference	84
VI. Skilled Memory, Mnemonics, and Levels of Processing	93
VII. Serial Outputs and Chunking	101
VIII. Workload and Working Memory	106
IX. Working Memory in Learning and Skill Acquisition	108
X. Final Comments	112
References	114

THE INTELLIGENT HAND

Roberta L. Klatzky and Susan J. Lederman

I. The Curious Discrepancy between Two Phenomena	122
II. Haptic Apprehension and Recognition: Theoretical Issues	128
III. Conclusions and Applications	147
References	149

SUCCESSIVE APPROXIMATIONS TO A MODEL OF HUMAN MOTOR PROGRAMMING

David A. Rosenbaum

I. Introduction	153
II. Hierarchical Decisions in Sequence Choices	154
III. The Motor-Program Editor Model	157
IV. Further Tests of the Hierarchical Decisions Model and Motor-Program Editor Model	163
V. The Hierarchical Editor Model	167
VI. Parallel Editing and Execution	175
VII. Conclusions	179
References	181

MODULAR ANALYSIS OF TIMING IN MOTOR SKILL

Steven W. Keele and Richard I. Ivry

I. Introduction	183
II. Issues in the Study of Timing	184
III. Individual Differences in Timing	189
IV. Further Analysis of Force Control and Maximum Rate	192
V. Individual Differences in Skill	197
VI. Neurological Analysis of Timing	203
VII. Other Approaches to Modularity	214
VIII. Conclusions	224
References	226

ASSOCIATIVE ACCOUNTS OF CAUSALITY JUDGMENT

David R. Shanks and Anthony Dickinson

I. Introduction	229
II. Contiguity and Contingency	230
III. Acquisition of Causality Judgments	236
IV. Blocking by the Causal Background	242

V. Retrospective Evaluation	247
VI. Comparator Theories	249
VII. Conclusion	256
References	258

ANXIETY AND THE AMYGDALA: PHARMACOLOGICAL AND ANATOMICAL ANALYSIS OF THE FEAR-POTENTIATED STARTLE PARADIGM

Michael Davis, Janice M. Hitchcock, and Jeffrey B. Rosen

I. Introduction	264
II. The Fear-Potentiated Startle Paradigm	264
III. The Pharmacology of Fear-Potentiated Startle	267
IV. Neural Systems Involved in Fear-Potentiated Startle	270
V. Sensitization of Startle by Footshocks	290
VI. Anxiety and the Amygdala	292
VII. Summary and Conclusions	296
References	297
<i>Index</i>	307
<i>Contents of Recent Volumes</i>	317

AN INTEGRATED COMPUTATIONAL MODEL OF STIMULUS-RESPONSE COMPATIBILITY AND PRACTICE*

Paul S. Rosenbloom

DEPARTMENTS OF COMPUTER SCIENCE AND PSYCHOLOGY
STANFORD UNIVERSITY
STANFORD, CALIFORNIA 94305

Allen Newell

DEPARTMENT OF COMPUTER SCIENCE
CARNEGIE-MELLON UNIVERSITY
PITTSBURGH, PENNSYLVANIA 15213

I. Performance: Compatibility and Goal Hierarchies	3
A. Data: Stimulus-Response Compatibility	4
B. Model: Goal Hierarchies	7
C. Results: Compatibility Simulations	16
II. Learning: Practice and Chunking	26
A. Data: Practice	27
B. Model: Chunking	29
C. Results: Practice Simulations	38
III. Discussion	43
References	49

Consider the position of a subject in a typical reaction-time experiment. On each trial of the experiment the subject is presented with a stimulus display—visual, auditory, or kinesthetic—containing some information from which he must determine the proper response, which is usually a vocalization or a manipulation by the hands or fingers. A single condition of the experiment defines a particular task environment—a set of possible stimuli (the stimulus environment), a set of possible responses (the response environment), and the mapping of the stimulus displays into the responses. The entire experiment, from the point of view of one within-condition subject, consists of a sequence of trials from this task

*This article is condensed from Rosenbloom, P. S. (1986), The chunking of goal hierarchies: A model of practice and stimulus-response compatibility. In J. E. Laird, P. S. Rosenbloom, & A. Newell (Eds.), *Universal subgoaling and chunking: The automatic generation and learning of goal hierarchies*. Hingham, ME: Kluwer Academic Publishers. Reprinted by permission of the publisher.

environment. This basic paradigm holds over a number of experimental domains, including stimulus-response compatibility and practice.

Theories of reaction-time phenomena usually focus on one particular domain, modeling well the data so circumscribed but effectively ignoring that each of these domains represents only a single aspect of what is in fact an integrated performance by the experimental subjects. When it comes time to build an integrated model of performance, this kind of approach lends itself best to a big-switch model. That is, each of the individual domain models is included as an independent component of the overall model, and a big conceptual switch selects the appropriate submodel when the experiment to be explained falls within its domain. The resulting model is little more than the sum of its parts and has special difficulties with situations that require interactions between the submodels.

An alternative approach to modeling reaction-time phenomena is to do it in the context of a cognitive architecture (Newell, 1973; Anderson, 1983; Pylyshyn, 1984; Laird, Newell, & Rosenbloom, 1987). A cognitive architecture specifies the set of fixed mechanisms upon which cognition is based. A complete architecture would specify the mechanism underlying learning, memory, problem solving, perception, motor behavior, etc. Human performance in a variety of domains can then be modeled as the interaction between the architecture and knowledge about tasks and strategies. Doing so has a number of potential benefits. First, it ensures that the various domain models can all coexist within a single working system. A big-switch model has problems when a task requires the interaction of two phenomena to be produced by submodels with incompatible assumptions. Second, architectures embody a set of mechanisms which may individually, or through their interaction, produce the desired phenomena without further assumptions. The model will be much more than the sum of its parts when a small set of basic mechanisms interact to produce all of the phenomena. Third, the architectural mechanisms are usually motivated by need—the system will not run adequately without them—while mechanisms hypothesized to explain reaction-time phenomena are usually motivated by the degree to which they match the data. If a mechanism meets both of these criteria, its likelihood of being correct is greatly increased. Fourth, the reaction-time phenomena appear as side effects of a working system actually trying to perform the experimental tasks, just as they do in psychology experiments. And fifth, these studies can be a good way of investigating the nature of the cognitive architecture itself.

In this article we present a theory of two reaction-time domains: stimulus-response compatibility and practice. This theory consists of two components: (1) a model of task performance, based on the concept of goal hierarchies, and (2) a model of learning, based on the concept of chunking. The compatibility and practice effects are produced by first constructing models of how subjects are performing specific experimental tasks, and then simulating these models to determine the time required to perform the tasks. The compatibility phenomena arise because of the differences

between the task-performance models underlying subject behavior in the different compatibility conditions. The practice phenomena arise because of changes wrought by the learning model to the task performance models. Though these two components are discussed independently in this chapter, they are actually two integral parts of a single system capable of both performance and learning. Learning occurs in stimulus-response compatibility situations, and it is impossible to run a practice experiment without having subjects actually perform some task.

The theory is implemented as a goal-based production-system architecture called Xaps3 (described in Rosenbloom, 1983; Rosenbloom & Newell, 1986). Though some of the architectural assumptions in Xaps3 are direct reflections of the theory, most of the assumptions are shared with the wider class of production-system architectures such as Act* (Anderson, 1983) and Ops5 (Forgy, 1981). The Xaps3 implementation of the theory has been used to generate simulated timing results for four experiments from the compatibility and practice literature (Seibel, 1963; Duncan, 1977; Fitts & Seger, 1953; Morin & Forrin, 1962).¹ The Seibel (1963) experiment has been instrumental in driving our theoretical work on practice (Newell & Rosenbloom, 1981; Rosenbloom & Newell, 1982; 1987). It is used to evaluate the practice predictions of the model. The Duncan (1977), Fitts and Seeger (1953), and Morin and Forrin (1962) experiments are three of the major stimulus-response compatibility experiments that are relatively free of confounding phenomena. They are used to evaluate the compatibility predictions of the model. In addition, the Duncan (1977) and Fitts and Seeger (1953) experiments provide practice data that can be used as a further evaluation of the practice model, and of the interaction between compatibility and practice.

The next three sections present the data to be modeled, the model, and the simulated results. This presentation is divided into a section on the performance model and one on the learning model. Stimulus-response compatibility is discussed in the performance section, and practice is discussed in the learning section. The final section contains a discussion of the model along with some potential objections to it.

I. Performance: Compatibility and Goal Hierarchies

In this section we lay out the basis for task performance, that is, how the actions of an experimental subject are determined by the interaction of the experimental stimuli with the subject's existing cognitive structures. This model, based on the concept of goal hierarchies, has been developed to

¹Because the model currently says little about the sources of performance errors, this work is focused entirely on timing data. For a speculative discussion of errors in this context, see Rosenbloom (1983).

model the main stimulus-response compatibility phenomena and to form the basis of performance in practice tasks. The body of this section consists of presentations of the relevant compatibility data, the model of performance, and results generated by the model.

A. DATA: STIMULUS-RESPONSE COMPATIBILITY

It was known by the early 1950s that the stimulus and response environments in an experimental situation could not be considered independently (Fitts & Seeger, 1953). The interaction between the two, as defined by the mapping, is often critical. This phenomenon is termed *stimulus-response compatibility*. Consider a concrete example in which there are two buttons, one above the other, that can be used to summon either an up elevator or a down elevator. In the *compatible* situation, the upper button summons the up elevator and the lower button summons the down elevator. In the *incompatible* situation, the relationship is reversed—the upper button summons the down elevator and the lower button summons the up elevator. In the compatible situation, people are faster and make fewer errors. These effects are robust and rather large. The problems encountered in performing in the incompatible situation do not stem from a lack of knowledge about the correct relationship—subjects learn the mapping from stimulus to response before the experiment begins—instead, it is a problem in actually performing the mapping.

Turning to the experimental work on compatibility, the most straightforward instances of the phenomena occur when the stimulus and response environments do not vary across conditions; only the mapping varies. In Duncan (1977), the stimulus environment consisted of an oscilloscope on which a vertical line could appear in one of four horizontal positions (top part of Fig. 1). The response environment consisted of a set of four buttons, lying under the fore- and middle fingers of the subject's hands (bottom part of the figure). On each trial of the task, one of the lines would appear on the oscilloscope and the subject was to press the appropriate button. There were three conditions in the experiment, each of which specified a different mapping of line position to button. In the corresponding condition (Fig. 1a) each line was mapped to the button below it. In the opposite condition (Fig. 1b) each line was mapped to the opposite button—the first line to the last button, the second line to the third button, the third line to the second button, and the last line to the first button. In the remaining mixed condition (Fig. 1c) half of the combinations (either the inner two or the outer two) were corresponding and the other half were opposite.²

²Duncan actually employed both mixed conditions, one in which the inner two lights were corresponding and one in which the outer two were. However, because we are not currently modeling differences in discriminability, we do not distinguish between these two variations.

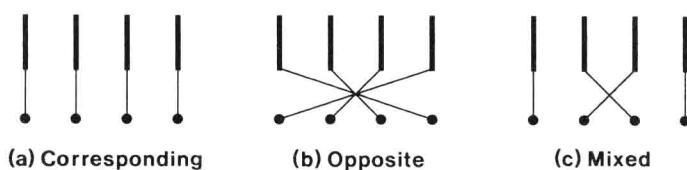


Fig. 1. The three compatibility conditions in Duncan (1977).

Table I shows the reaction times for the three conditions. Though the mixed condition is only a single condition, each trial is itself either corresponding or opposite. Therefore, the data have been partitioned to reflect this split. The main thing to notice at the moment is that though the stimulus and response environments remain unchanged, manipulating the mapping yields differences in the time it takes to perform the task. The opposite trials are consistently slower than the corresponding trials, and the trials in the mixed condition are slower than the analogous ones in the nonmixed (pure) conditions. In fact, the two factors appear to be independent, with an extra 60 msec for the opposite mapping, and 100 msec for a mixed condition.


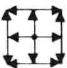


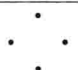

Even when the stimulus and response environments are modified across conditions, it is the mapping that accounts for most of the variance. Fitts and Seeger (1953) reported a nine-condition experiment in which three stimulus environments were crossed with three response environments. Table II shows the apparatuses used to define these environments as well as the mean reaction times for the nine conditions. Stimulus apparatus S_A contains eight lights at the 45° points of a circle. On each trial in which it is used, exactly one light goes on. Stimulus apparatuses S_B and S_C both contain four lights. On each trial either one light comes on or a pair of adjacent lights (with respect to the circle) comes on. With apparatus S_B the four lights are at the 90° points of a circle. With apparatus S_C the display is twice as wide, the horizontal lights in one half, and the vertical lights in the other. Adjacency for apparatus S_C is defined as if the lights were still in the circle of apparatus S_B . The light on the far left is “at” -90° , the middle light is “at” 90° , the top-right light is “at” 0° , and the bottom-right light is “at” 180° .

TABLE I

MEAN REACTION TIMES (IN MSEC) AND MARGINAL DIFFERENCES FOR THE FOUR TYPES OF TRIALS IN DUNCAN (1977)

	Corresponding	Opposite	Δ
Pure	431	489	58
Mixed	529	590	61
Δ	98	101	

TABLE II
MEAN REACTION TIMES (IN MSEC) FOR THE NINE
CONDITIONS IN FITTS AND SEEGER (1953)

	R_A 	R_B 	R_C 
S_A 	390	430	580
S_B 	450	410	580
S_C 	770	580	480

The three response apparatuses are defined analogously to the stimulus ones. In response apparatus R_A there is a lever that can be pushed toward any of the eight 45° angles. When used in conjunction with S_A , the lever is pushed in the direction of the light. With S_B and S_C , if one light is on, the lever is pushed in that direction; if two lights are on, then the lever is pushed toward the mean of the two angles. For example, if the right and top lights (which actually appear in the middle and top-right of the display, respectively) in apparatus S_C are on, then the lever should be pushed at a 45° angle. Response apparatus R_B allows the lever to be pushed at only 90° angles. When it is used with either stimulus apparatus S_B or S_C , the lever is pushed in each direction specified by an on light. This may require the lever to be pushed either once or twice. When used with stimulus array S_A the lever is pushed once if the light is at a multiple of 90° and twice otherwise (at an angular displacement of $+45^\circ$ and -45° from the light that is on). Response apparatus R_C is analogous to R_B except that it requires two hands to manipulate, one for each of the two orthogonal directions. For all three response apparatuses the reaction time is measured up until the first movement is begun. Because movement time is not included, two movements need not take longer than one.

The first thing to notice about the reaction times for these conditions is that in each row and column in the table the fastest reaction time belongs to the one on the main diagonal. For each stimulus apparatus there is a different response apparatus that produces the fastest times. In the analysis of variance reported by Fitts and Seeger, the effects of stimulus apparatus and response apparatus individually were significant, but in their words, "The variance that can be attributed to interaction is very much larger than the variance attributable to the primary effects of either stimulus or response sets alone" (p. 204).

This experiment also reveals that just labeling conditions as compatible and incompatible can miss a good deal of what is going on. Though the conditions on the main diagonal are compatible and those off the diagonal are incompatible, some of the incompatible times are faster than some of the compatible times. A theory of stimulus-response compatibility should explain these complexities.

Both the Duncan and the Fitts and Seeger tasks are spatial tasks, but the phenomena of compatibility are not so limited. Morin and Forrin (1962) used a set of five symbolic tasks (Table III). In Conditions I and IV, each trial consisted of the visual presentation of an arabic numeral to which the subject was to respond by saying the number (for example, see "2", say *two*). The conditions differed in the number of alternatives in the task environment (2 and 4, respectively). In Conditions II and V, each trial consisted of the visual presentation of a symbol (+, ■, ●, or ▲) to which the subject was to respond by saying a number that had been associated with it (4, 7, 2, and 8, respectively). Again the conditions differed in the number of alternatives that could be presented. Condition III was a mixed condition in which the stimulus environment consisted of two numbers (2 and 8) and two symbols (+ and ■). In Table III, Condition III has been split according to whether a number (IIIa) or a symbol (IIIb) appeared as a stimulus on the trial.

The reaction times for these conditions divide up into three groups separated by about 100 msec each. Conditions I and IV are the fastest (the "compatible" conditions), at around 500 msec. At around 600 msec we find Conditions II and IIIa, and at around 700 msec we find Conditions IIIb and V.

B. MODEL: GOAL HIERARCHIES

The performance model is based on work in three different areas: (1) stimulus-response compatibility theory, (2) applied information-processing psychology, and (3) artificial intelligence.

TABLE III
COMPATIBILITY CONDITIONS AND RT (IN MSEC) FROM
MORIN AND FORRIN (1962)

Condition	S-R Pairs		RT	Condition	S-R Pairs		RT
I	2-2		520	II	+ -4		590
	8-8				■ -7		
IIIa	2-2	+ -4	600	IIIb	+ -4	2-2	710
	8-8	■ -7			■ -7	8-8	
IV	2-2	4-4	490	V	+ -4	● -2	720
	8-8	7-7			■ -7	▲ -8	

1. Stimulus-Response-Compatibility Theory

Though the phenomena of compatibility are important—especially in human factors, where the relationship between stimulus displays and simple manual controls is critical—and have been studied since the early 1950s, there are still no useful theories of compatibility. Welford (1980) sums up this situation nicely: “Surprisingly, after more than twenty years’ research there is no metric for compatibility that transcends particular experimental conditions” (p. 99). Despite the lack of any complete theories, some theoretical statements have been made about aspects of compatibility. The earliest and most widely accepted is that compatibility is a function of the transformations or encodings that must be performed on the stimulus to yield the response. Deininger and Fitts (1955) put it this way: “The fundamental assumption for the model used in developing the S-R compatibility concept is that all perceptual-motor performance involves one or more steps of information transformation, since the response must always be encoded in a manner different from the stimulus” (p. 318). In addition, they went on to propose one other factor they felt was important—whether “the pairings of stimulus and response elements agree with strong population stereotypes.”

Brebner (1973) took the Deininger and Fitts model one step further, by considering what happens in a multistage transformation between stimulus and response. He proposed that in more compatible situations a single recoding or translation process can be applied to all of the steps. For example, Brebner’s hypothesis would imply that it is easier to do a string of three additions than two additions with a subtraction sandwiched between them.

The closest anyone has come to a metric was in the work of Morin and Grant (1955), though it was still for only one set of task variations. They examined a set of tasks based on eight lights and eight response keys, in which the mapping of light to key could be changed. On each trial, a pattern of lights was presented and the subject pressed the appropriate keys. Morin and Grant compared reaction time with the rank correlation of the mapping between the stimuli and response (that is, the serial positions of the lights were compared with those of the buttons associated with them) and found an effect of the absolute value of the correlation (a quadratic effect) and the sign (a linear effect). Shepard (1961) showed that a partial account for this effect can be obtained by looking at the stimulus and response generalizations (or confusions) and the permutation matrix defining the mapping between stimuli and responses.

To explain the results of mixed conditions, such as in the Duncan (1977) study described earlier, Smith (1977) proposed a parallel-iterative scheme. In this theory, each S-R pair has an association time, reflecting some unspecified function of the compatibility of that pair. Each iteration of the

process takes time proportional to a weighted sum of the association times for all of the pairs in the condition (not just of the pair for the current trial). Across iterations, excitation of the responses accumulates until one response reaches its threshold. Total reaction time is therefore the sum of the times required for all of the iterations until a response. This theory predicts that the time to do a mixed mapping will be between the times to do the respective pure mappings, because the corresponding mapping consists of four fast associations, the mixed mapping of two fast and two slow associations, and the opposite mapping of four slow associations. It also predicts that the addition of more S-R pairs to the condition will always raise the reaction time for the other pairs (though the effect may be small if the association times for the new pairs are small).

Smith's theory treated each S-R pair as a distinct component in the response selection process. In contrast, Duncan (1977) proposed that "responses may be selected, not on the basis of individual S-R associations, but by use of a rule or system of rules." Duncan concluded that there were two factors determining spatial stimulus-response compatibility:

Spatial CRT [Choice Reaction Time] is influenced by two different properties of the mapping. One is the spatial relationship of individual S-R pairs, this is, in terms of the proposed model, which transformation must be made. The other is the number of different relationships in the whole mapping; that is, the number of transformations from which selection must be made. (p. 60)

2. *Applied Information-Processing Psychology*

Card, Moran, and Newell (1980, 1983) proposed that, for tasks involving cognitive skill, a procedural representation of task performance provides good approximations to human performance. They have worked this out for one domain, computer text-editing and text-based command language interaction. They showed that models based on the concepts of goals, operators, methods, and selection rules—GOMS models—provide an excellent account of the behavior of human subjects.

When this idea is paired with the earlier theoretical ideas, it yields the image of performance in compatibility tasks as being mediated by procedures or algorithms. An algorithm for a condition in a stimulus-response compatibility experiment specifies a sequence of operations that will compute the correct response for any stimulus that the subject may face during that condition. This idea is the basis for the *GOMS model of stimulus-response compatibility*.³ Compatibility phenomena are produced because the more incompatible tasks require algorithms that take longer to perform. In

³This model was referred to as the *algorithmic model of stimulus-response compatibility* in Rosenbloom (1983). The name has been changed to emphasize its relationship to the GOMS framework.

Duncan's terms, this is either because the "rule" for the condition is complex or because there are many rules for the condition (a mixed condition), necessitating numerous tests and branches in the algorithm.

Given the specification of the GOMS model, the path to a compatibility metric is straightforward. The first step is to analyze the task situations to ascertain the algorithms employed in their performance. There may be one or more such algorithms, reflecting different strategies being employed by different subjects, or the same subject at different times (see Newell, 1973; Baron, 1978, for discussions of the issues surrounding the availability of multiple methods). One approach to the development of task algorithms is to perform a set of experiments geared toward ascertaining the algorithms actually used by subjects in the various task situations. Another approach is to perform an abstract task analysis similar in nature to the process a programmer goes through in developing an algorithm for a task. In this article we use a form of this abstract task-analysis approach exclusively. However, to increase the likelihood of developing algorithms that reflect what human subjects actually do, we have imposed two biases on this algorithm-development process. The first, and most important, bias is to not include mechanisms in the algorithms that violate in significant ways what is known about the capabilities and limitations of the human cognitive architecture. The second bias is to assume that subjects will tend to use the simpler (faster) algorithms when many are possible.

Once algorithms have been derived, the second step is to perform a complexity analysis of the algorithms (Aho, Hopcroft, & Ullman, 1974). This involves assigning cost measures (i.e., amount of time to perform) to the primitive steps in the algorithm and determining how many of each type of step would be executed in the process of performing a trial. Given these values it is relatively straightforward to predict the mean reaction times for the conditions of the experiment. For each algorithm that might be used in a particular condition, the mean cost of performing a trial in that condition is determined. If there is more than one algorithm for the condition, then the costs from the different alternatives must be merged into a single cost for the condition. Since we have no data on the relative frequencies of the alternatives, we make the minimal (Laplacian) mixing assumption: assume the algorithms are equally likely (across subjects) and take their mean behavior.

We have elsewhere presented an approximate version of the GOMS model of compatibility that facilitates quick back-of-the-envelope style calculations for predicting relative reaction times between conditions in compatibility experiments (Rosenbloom, 1983) and shown how such a

model can be useful in the domain of human-computer interaction (John, Rosenbloom, & Newell, 1985; John & Newell, 1987). In this article we present a goal-hierarchy version of this model that, while more complex than the GOMS model, fits in with current notions about cognitive architecture and facilitates its integration with the model of learning.

3. *Artificial Intelligence*

Goal hierarchies are a common control structure for the kinds of complex problem-solving systems found in artificial intelligence (Rich, 1983), but this is the first time they have been applied to the domain of reaction-time tasks. The foundation concept of the goal hierarchy is that of the goal—a data structure representing a desired state of affairs. A goal is not a procedure for bringing about the desired state; it is only a description of the state. In order to bring about the goal state, there must be one or more methods associated with the goal. A method could be a rigid algorithm, or it could be one of the more flexible weak methods (Newell, 1969) such as means-ends analysis (Ernst & Newell, 1969) or heuristic search (Nilsson, 1971). The current model is overly simplistic in that it conflates the properly distinct notions of goal and method into a single active concept. These conflated “goals” are active processes, much like functions, that take a set of parameters and return a set of results. This simplified treatment is sufficient for the reaction-time experiments modeled here because these tasks require very little in the way of sophisticated problem solving. More complex tasks require a more sophisticated model (see Laird, 1983; Laird *et al.*, 1987, for one such model, which is also closely related to the current work).

A single goal generates a goal hierarchy when the goal can be decomposed into a set of simpler goals, and those goals can be decomposed even further. The recursion terminates when the goals are so simple that they can be attained directly. To a large extent the terminal goals in the hierarchy do the actual work. Nonterminal (or internal) goals create the control structure that organizes the performance. Figure 2 shows the goal hierarchy we developed for the Duncan corresponding task. The top half shows the bare bones of the hierarchy as a tree structure. The nodes are labeled with a number representing the order in which the goals are processed in a depth-first fashion. In depth-first processing there is always exactly one goal being actively worked on at any point in time. We refer to this goal as the active or current goal. When a goal generates a subgoal, the subgoal becomes the active goal, and the parent goal is suspended until control is returned to