Note that each value of y_i may be characterized by a different number of points, J_i, depending on the record, $\mathbf{H}_i$. In choosing a decision criterion we must deal with the problem that K, the number of points in the current input, $\mathbf{X}$, is not necessarily equal to J_i, for $i = 1, 2, \ldots, M$. The rule used in our model is based on the number of points in X as shown in the following:

$$\max_{i = 1,2 \ldots M} i \left[\prod_{k=1}^{K} \left\{ \sum_{j=1}^{J_i} \frac{1}{J_i} \left(P(\mathbf{Y}_{ij} = \mathbf{X}_k | \mathbf{H}_{ij})\right) \right\} \right]$$

The probability density at each point designated by subscript k of X_k is averaged over all points J_i for each record. This probability density function has N equally distributed Gaussian components of variance, σ^2, and is thus given by

$$P(\mathbf{Y}_{ij} = \mathbf{X}_k | \mathbf{H}_{ij}) = \frac{e^{-\sum_{n=1}^{N} \frac{(X_{kn} - H_{ijn})^2}{2\sigma^2}}}{(\sqrt{2\pi\sigma^2})^N}$$

Finally, for the case in which more than one record has the same response label, *max i*, in (3) is replaced by *max c*, and the term between the outermost brackets is summed over the number of records that have the same response label, yielding an expression analogous to (2).

7 Intelligent Pattern Recognition: Hierarchical Organization of Concepts

Michael L. Commons
Peter W. Hallinan
(with Wilson Fong and Katherine G. McCarthy)
Harvard University

INTRODUCTION

This chapter describes the properties of intelligent pattern recognizers and the failure of Artificial Intelligence theory to attend to these properties. Towards the end of the nineteenth and early in the twentieth century two profound insights into the nature of the world came about. Both attack the ideal rationalist explanation of thought and perception. First, in 1929 Szilard (cited in Resnikoff, 1985) realized that decreasing entropy[1] corresponds to a gain in information[2]. Second, Spencer (1870), Jennings (1906), and Thorndike (1898, 1911) established that an organism's experiences and preferences determine its behavior.[3] Spencer articulated the idea that levels of intelligence could be viewed in terms of levels of correspondence between internal states and external circumstances; Jennings, as a significant contributor to the evolving science of behavior qua behavior; Thorndike, for that kind of contribution as well, but also for the law of effect, a principle that helps explain how internal state and external circumstance are coordinated. Coupled together, these insights provide a basis for arguing that some machines[4] will order their internal states[5] in response to machine–environ-

[1] Entropy is a measure of the disorder of a system. For example, a material that changes it state from liquid to gas increases its entropy.

[2] An observer gains information about a physical quantity such as mass by increasing the resolution of his measurement of the quantity.

[3] Whereas Hobbes (1650, 1651) said that experiences and preferences determine behavior, empirical support came much later.

[4] "Machine" refers to any computational device, whether biological or manmade.

[5] Internal states of the brain represent the tendency to respond.

ment interactions. These interactions will constrain the values that the internal state may assume and will affect the state of the environment. A machine's intelligence increases with the complexity of its internal order and the machine-environment interaction. Intelligence, therefore, is based on the notion of adaptive information acquisition and preservation.

In contrast, Artificial Intelligence (AI) theory assumes that understanding intelligence depends on understanding knowledge representation and manipulation. Knowledge acquisition is of secondary importance. AI has had only moderate success in developing machines that function usefully in real domains. Unless AI theorizing includes models of adaptive mechanisms, AI machines will always depend on human beings for programming, reproduction, and evolution and thus never function as independent entities capable of solving a wide range of problems.

True intelligence requires, among other things, the ability to recognize patterns. Ascertaining conditions necessary for intelligent pattern recognition can help illuminate the shortcomings of current AI theory and suggest a truly intelligent cognitive-adaptive mechanism. This analysis relies on evidence from biology and stage change theory. We argue that current AI theory with its task orientation fragments the nature of intelligence, which is not arbitrary in the organic world. The construction of an intelligent machine depends not simply on designing systems that accumulate data and apply decision rules for manipulating that information, but on certain organizing principles that are inferable from the natural evolutionary process. These biological and psychological principles are the focal points of this analysis.

REQUIREMENTS FOR INTELLIGENT PATTERN RECOGNITION

Any intelligent pattern recognition machine (IPR) must be able to isolate causal relations. Otherwise the machine cannot know how to change its behavior to compensate for changes in the availability of such things as food, mates, and danger. To isolate causal relations, a machine must possess the requirements listed in Table 7.1. Although these conditions may not suffice to characterize intelligent pattern recognition, they do correlate well with observed animal behavior (Dore & Dumas, 1987; Herrnstein, in press) and human behavior (Commons, Richards, & Armon (1984).[6] We first discuss the nature of these principles in detail and then present results demonstrating the necessity of all conditions.

[6]We do not make the distinction that humans are intelligent and animals are not. Instead we say that humans possess greater internal complexity and thus greater intelligence.

TABLE 7.1
Isolating a Causal Relation

A machine must
(1) exist in a regular, complex, and dynamic environment
(2) interact with its environment using input and output devices
(3) display a preference for certain internal states
(4) adapt actively using feedback
(5) store past interactions with its environment
(6) organize its concept formation hierarchically

The Environment: Regular, Complex, and Dynamic

The environment containing the IPR must have three characteristics. First, its structure must be regular, i.e., predicted on the lawfulness of cause and effect. It must be possible for the machine to construct a model of the environment that can be at least partially successful in predicting future states of the environment. Second, the environment must be complex, i.e., nonuniform and multivariate, with regard to its own structure, its inputs to and outputs from the machine. The complexity ensures that tasks posed to the machine by the environment cannot be solved by simple search or guessing. Third, variables representing the environment state, machine inputs, and machine outputs must change with time so that the machine cannot solve all its tasks by converging to a fixed set of behaviors. In a static environment, having intelligence is not advantageous.

The regularity, complexity, and dynamism are interdependent; without complexity and dynamism there is no need for sophisticated processing, and without regularity there is no possibility of processing. The more numerous the degrees of freedom and the greater the nonuniformity involved in the environment structure and machine inputs, the greater the internal order of the machine must be to succeed. A measure of the complexity of external inputs and structure that must be accommodated by a successful machine is thus a measure of the intelligence of the machine.

Interaction with the Environment

The prevailing tendency of physical systems is towards increasing disorder, or entropy, which corresponds to a loss of information. Thus the environment in which the machine is embedded can be considered an entropy maximizer. The machine, however, continuously expends energy both to increase the scope, abstraction, and accuracy of its internal representation of its environment, and to preserve via reproduction at least some of the internal representations it has constructed. Whenever the machine updates or generalizes its internal representation, it gains information and decreases entropy at its locale within the environment.

Thus, as an entropy reducer, the machine stands in continuous conflict with its environment, an entropy maximizer. Moreover, the conflict is unequal because the machine is by definition possessed of less powerful resources than the environment; the machine is less complex and occupies just one location within the environment. Therefore, the interaction is environment driven, or input driven. The environment will pose challenges to the machine, and not vice versa. In the natural world, these challenges are posed via competition for scarce resources, e.g., mates and food, and the hostile nature of changes in the environment, e.g., a change from warm to cold temperature. The response of the machine to input is an action on the environment that must have local effects persistent enough to influence future interactions. In summary, the system can be described as a state machine, with events at time (t − 1) influencing events at time (t). Of course, the state at time (t + 1) cannot affect behavior at time (t).

Additionally, all interactions must be predictable, though not necessarily deterministic. The dynamism principle, which prevents the machine's repertoire of behaviors from converging to a fixed set, also prevents interactions from reaching equilibria, rather than forcing interactions to be unmodelable. Even novel inputs must have predictable results, though novelty may eventually alter the behavioral repertoire. If the environment structure represents at least some of the physical laws (e.g., predator–prey relations, fluid turbulence) of our own world, then the need for regularity constrains parameter values in nonlinear models of machine–environment interaction such that relations of interest to the machine cannot be instantiated with values leading to chaotic behavior. Although intelligence obviously cannot develop in stable situations, it also cannot develop out of interactions describable by chaotic nonlinear feedback models, as in the Verhulst (Peitgen & Richter, 1986) model for chaotic population growth.

Preference for Certain Internal States

Entropy minimization occurs because the machine has a preference for certain internal states. Preference can be understood as a mechanism that evaluates the effect of an action by the environment on the organism. The machine can then use the evaluation to change the tendency of a behavior in response to the interaction. In the organic world, preference can be interpreted as a cell's need for oxygen, nutrients, neurotransmitters, or other hardwired biochemical needs. There may be other internal preferences. Without internal preference, success is undefined, reinforcement cannot affect behavior, and intelligence cannot develop. The existence of preference implies that there must be a cost for deviating from preference.

Preference provides a mechanism for selecting among goals. Preference does not imply that the organism works towards a goal such as a reward, however. Changes in the environment dictate the future changes in the machine's behavior. For example, people are more willing to work for the promise of a reward after previous promises have resulted in reward. They do not work for the reward itself, which has not occurred, but for the promise.

Therefore, the existence of preferences and goals does not imply that the machine changes state in a teleological fashion. Rather, they constrain the set of possible responses that the machine might make to a given internal or external stimulus. The criteria for choosing a response at any level of abstraction, whether the burn reflex or language, does not equate to having a goal. Goals and plans are internal concepts that are formulated as a response to internal or external stimuli (Skinner, 1969) and therefore are nonteleological in nature despite their use in constraining the machine's behavior. Plans and goals are explicit statements of rules and outcomes, whereas preference may describe an implicit relation between behavior and outcomes. At the level of hardwired primitives of the machine, preference does equate to functional output, but at cognitive levels preference equates to bottom-up[7] constraints on stimulus–response pairings.

Active Adaptation Using Feedback

Novelty is implicit in the nonuniformity of the environment and the localization of the machine. The existence of novelty requires that the machine adapt its behavior to satisfy internal preferences. Adaptation is not instantaneous; it requires a feedback loop in which (a) the environment acts on the machine, (b) the machine, somehow modified by the interaction, acts on the environment, (c) the environment acts on the machine again, and so on. In this process, the machine responds according to the relative success or failure of each of its successive actions; that is, it learns through reinforcement.

Vaughan and Herrnstein (1987) have described two kinds of adaptability. The first is intragenerational—affecting just one machine. There are three subtypes: (1) short-term changes in response—e.g., ducking a snowball; (2) long-term changes in response—e.g., learning a language; (3) metachanges—changes in responses to their own responses—e.g., changing stage.

There are two classes of intragenerational learning mechanisms: supervised and unsupervised. Supervised intragenerational learning depends on the programmer to make changes in the programs. Unsupervised intragenerational learning requires no other machine to modify its programs. Both depend on the IPR and contrast sharply with all other methods of knowledge acquisition, which are independent of the IPR, e.g., programming a computer.

The second kind of adaptability is intergenerational—affecting machines across generations. The assumption of a dynamic environment implies that the frequency with which a response is used will vary. If those responses that prove very useful are hardwired, they will not have to be relearned by each machine. Because usefulness is determined by output, the system is self-stabilizing rather than self-optimizing (Vaughan & Herrnstein, 1987). More generally, as Vaughan and Herrnstein suggest, natural selection is an Evolutionarily Stable Strategy

[7] Bottom-up processing starts from the input and works upward.

(ESS), not an optimizing strategy. In biological terms, hardwiring very useful knowledge corresponds to favoring those mutations that hardwire such knowledge. In an artificial IPR, we remove the randomness from evolution, without crippling natural selection.[8]

Adaptive systems of any kind require both a measure of how much they must adapt and a mechanism to accomplish the adaptation. These measures and mechanisms occur at many different levels. At the highest level, the measure is the continued existence of the IPR. At a lower level, e.g., removing a hand from a flame, the measure may be the intensity and duration of signals along the afferent pathways from the hand. At any level, though, the mechanism driving adaption can be schematized as a feedback pathway. AI models without feedback will fail to adapt. Hebb (1949) has described the essential component of biological feedback as the adjustment of cell connections in proportion to their usage. Feedback has other implementations, though. It can be an efferent copy, a negative image of expected or self-generated sensory input sent to the lower levels in the hierarchy to mask such input (Bell, 1981). For example, electric fish must mask the outputs of their electroreceptors when triggered by their own electric organ discharges. Another kind of efferent copy is the presynaptic inhibition as used, for example, by the crayfish to suppress its tail-escape reflex that is stimulated by normal swimming motions.

The matching law (Herrnstein, 1970) and melioration (Herrnstein, 1982; Herrnstein & Vaughan, 1980) can be understood as the basis for self-programmability on an implementation level. In a situation, each activity engaged in may be represented by a program. Reallocating the proportion of time engaged in an activity in a way that tends to equalize the reinforcement value obtained from those allocations is a form of self-programming.

Using Hebb's (1949) connectivity notion metaphorically, melioration within a specific machine or organism might be thought of as synapse facilitation when information pathways that are reinforced are used more and more often. The direction of growth and the length of axons are genetically coded, implying that the number of neurons with which a given neuron can synapse is severely limited by DNA. Because learning manifests itself physically as the synapses formed and reformed (plasticity) by the neurons, certain configurations of networks will be evolutionarily favored.

Reproduction is not a preference of living organisms but an adaptive mechanism for preserving and collecting information over a period of time exceeding the life-span of a single machine or organism. Reproduction allows changes in the hardwiring and relearning of lower stage concepts for new environments.

Adaptation is constrained by the physical form of the machine, and by the cost of facing the novel hostilities that are necessary to produce adaptive generalization.

[8]New programs and machines replace old ones, the newer ones being based on the older.

Remembering Machine–environment Interactions

The need to remember or store the effects of the relationship between past events depends on the degree to which inputs are uniform and on the cost of failure. If stimuli never repeated, then remembering them or their effects would be useless. Learning would also be impossible. If failure has no cost, then intelligence has no advantage over stupidity. More precisely, the machine must respond to non-novel stimuli in a way that reflects its past success or failure in responding to the stimuli. Successes may be defined as outputs that meliorate hostile novelties or obtain preferred outcomes (Herrnstein & Vaughan, 1980), whereas failures may be described as the nonoccurrence or loss of reinforcement. To have successfully adapted, the machine must produce successes without having to learn each time it is presented with a similar situation.

Intelligent machines, however, remember in ways different from computers. Computers store specific data in localized physical memory cells. Instead, intelligent machines remember by changing tendencies to produce a particular response given a particular stimuli. (McClelland & Rumelhart, 1985; Rumelhart & McClelland, 1985) The machine as a whole constructs a representation of an event, and later it may construct a representation of a stream of events. But no individual element within the machine contains a memory of an event, even though the state of any element is influenced by the dynamics of the net. Changes in the network will produce changes in a given element. One cell does not define the action of the net even for a moment.

Hierarchical Organization of Processing

As Resnikoff (1985) has shown, hierarchy is a powerful organizational principle that unifies information processing in a variety of fields: neurophysiology, psychophysics, computer science, developmental psychology, etc. Examples of hierarchically organized neural nets abound across all phyla, sometimes with quite unambiguous functional explanations. Stage theory, as developed by Piaget (1952) and extended by Commons and Richards (1984a,b) and others (Campbell & Bickhard, 1986; Case, 1985; Fischer, 1980; Pascual-Leone, 1980), embeds human cognitive development from infant to adult in a hierarchically organized framework.

The need for hierarchical processing stems from the machine's need to respond to a wider and wider class of actions by the environment (Kehoe, in press). For example, in a mathematics domain, the machine must learn to deal with number, then variables, then functions, etc.

Hierarchical processing must meet four conditions. First, later actions are always defined in terms of earlier actions. Actions higher in the hierarchy depend on lower actions. Second, these higher actions organize the subsequent outputs from lower actions in the hierarchy. Third, this organization is not arbitrary.

Finally, the outputs of these organizing actions serve as inputs for the next, higher, stage in the hierarchy. Not all processing is hierarchical. Transfer of behavior across stimulus domains does not require construction of a new hierarchy but rather involves iteratively applying a set of actions from the same stage (Commons, Grotzer, & Davidson, in preparation).

Several clear examples exist of conceptual hierarchies that directly correspond to anatomical hierarchies. Newman and Hartline (1981) have demonstrated that in the rattlesnake optic tectum inputs from afferent neurons sensitive to visible and infrared light are integrated at the next level by interneurons that fall into six categories: neurons responding strongly to infrared and visual stimuli (I & V), to visual stimuli only (I-depressed V), to infrared or visual stimuli (I or V), etc. These second-level concepts are integrated to give more abstract concepts. Thermopositive visual objects (e.g., a mouse or a hot rock) will be perceived when the infrared and visual (I & V) interneurons respond strongly and visual-only interneurons respond weakly. Uninteresting (to the snake) thermoneutral visible objects will be perceived when the visual-only interneurons fire intensely.

A second example found by Matsubara (1981) also demonstrates how an additional level of problem solving contributes to more successful behavior. Matsubara has compared the electrolocation systems of the weakly electric fish Eigenmannia and Sternopygus. Electroreceptors of a fish respond to increases in amplitude of the fish's electric organ discharges caused by other fish or objects. Unfortunately for Eigenmannia, the electric fields of other fish can jam its electroreceptors. However, jamming does not happen to Sternopygus. The difference is in the electroreceptors. Eigenmania has two kinds: Type I, which is excited by amplitude increases, and Type II, which is inhibited by amplitude increases. In Sternopygus, though, 80% of its receptors are of Type III, a cell having two nonoverlapping receptive zones, one functioning as a Type I cell, and the other zone functioning as a Type II cell.

In Eigenmannia, jamming, or any large-field disturbance, triggers every Type I and Type II cell and generates noise (useless information) in higher processing centers. In Sternopygus though, a Type III cell will effectively sum the responses of each of its 2 zones and thus not respond at all. Only in the presence of a local, or small field, disturbance large enough to excite only the inhibitory zone or the excitatory zone will the Type III cell respond, and then it will respond as the cell type corresponding to the triggered zone would.

Thus, the jamming that works on Eigenmannia causes Sternopygus no trouble because it filters out large-field disturbances. Further, the filtering is accomplished by adding an intermediate level in the processing hierarchy.

HIERARCHICAL PROCESSING AND STAGE CHANGE THEORY

In the organic world, hierarchical processing may be observed in the sequential process of causality recognition. Commons, Grotzer, and Davidson (in preparation) studied the development of such processing in terms of concrete-to-formal-operational stage transition in 172 lower and middle-class fifth- and sixth-grade students over a period of 6 years, using one of eight versions of Commons, Miller, and Kuhn's (1982) laundry problem.[9] An examination of this experiment illuminates the hierarchical nature of learning. The mechanisms that were shown to be necessary for accomplishing stage change in this experiment can also suggest models for facilitating intelligent pattern recognition in AI theory.

The Laundry Problem: Introduction

In the Commons and Davidson (Davidson, 1983; Davidson & Commons, 1983) variation of the laundry problem, subjects predict how a cloth will come out when it has been washed in a given set of ingredients. The information given to subjects is dominated by a causal variable that can be isolated and identified by matching it with the outcome of the various episodes while at the same time disregarding all other variables. A detailed description of the laundry problem follows.

The laundry problems is composed of 16 episodes, each including 4 independent variables (soap type, water temperature, bleach brand, and booster color), followed by an outcome variable (cloth cleanliness) as shown in Table 7.2. The first 6 of these episodes are referred to as Informational Episodes. They give subjects enough information to determine the causal variable for the cloth outcome. Labeled bottles containing the actual cleaning agents were used to present the episodes. Each episode showed 4 combinations of pairs of washing agents: water, soap, booster, and bleach. Three of the Informational Episodes resulted in a cloth outcome of dirty; 3 in a cloth outcome of clean. The outcome was demonstrated by actual clean and dirty cloths. Only 1 ingredient made a difference; the outcome followed the others in a constrained, random way.

Procedure

Ten Test Episodes, presented one at a time, followed the six Informational Episodes. For each Test Episode, subjects were shown the four agents with which a stained cloth was washed. The Informational Episodes were displayed throughout the session so that subjects could refer back to them and so that the isolation of variables task was not confounded with a memory task. Subjects were asked to use the information from the Informational Episodes to determine the cloth outcome in each of the Test Episodes. The formal-operational task-required action was to detect which of four pairs of ingredients caused the cloth to come out clean.

[9] This problem was derived from Kuhn and Brannock's (1977) plant problem, which was derived from Linn, Chen, and Thier's (1976, 1977) plant problem, which was based on Inhelder and Piaget's (1958) pendulum problem (also see Siegler, Liebert, & Liebert, 1973).

TABLE 7.2

The Cloth was Stained with Red Lipstick. It Was Washed in Each of These 6 Ways

A Bleach	Powder Soap	Blue Booster	Cold Water	→	Dirty
B Bleach	Liquid Soap	Pink Booster	Hot Water	→	Clean
A Bleach	Powder Soap	Pink Booster	Hot Water	→	Dirty
B Bleach	Powder Soap	Pink Booster	Cold Water	→	Dirty
A Bleach	Liquid Soap	Blue Booster	Hot Water	→	Clean
B Bleach	Liquid Soap	Blue Booster	Cold Water	→	Clean

Look Back at the Examples. Now, Mark the Correct Ending

B Bleach	Powder Soap	Blue Booster	Hot Water	→	Clean ☐ Dirty ☐
A Bleach	Liquid Soap	Blue Booster	Cold Water	→	Clean ☐ Dirty ☐
A Bleach	Powder Soap	Pink Booster	Cold Water	→	Clean ☐ Dirty ☐
B Bleach	Liquid Soap	Blue Booster	Hot Water	→	Clean ☐ Dirty ☐
B Bleach	Powder Soap	Blue Booster	Cold Water	→	Clean ☐ Dirty ☐
B Bleach	Powder Soap	Pink Booster	Hot Water	→	Clean ☐ Dirty ☐
A Bleach	Liquid Soap	Pink Booster	Hot Water	→	Clean ☐ Dirty ☐
A Bleach	Powder Soap	Blue Booster	Hot Water	→	Clean ☐ Dirty ☐
B Bleach	Liquid Soap	Pink Booster	Cold Water	→	Clean ☐ Dirty ☐
A Bleach	Liquid Soap	Pink Booster	Cold Water	→	Clean ☐ Dirty ☐

Problem Presentation

Subjects were exposed to 16 problem presentations over approximately 5 months. The time between each test session varied. For each Test Episode an experimenter asked which outcome the subject thought would occur if the cloth were washed in the given combination of cleaning agents. The experimenter would then probe for the subject's descriptions and explanations. Subjects were asked which of all possible combinations of variables (i.e., soap; or soap and bleach; or soap, bleach, and booster; etc.) was most responsible for how the cloth came out, and whether or not one variable made a difference in the outcome.

In this version of the experiment, one group of subjects were told whether their answer was correct or incorrect (feedback), which they did not care about. Another group of subjects received verbal feedback and points (reinforcing feedback) for correct responses. These points were used in team competition for prizes. A subject received one point for hits (correctly predicting a clean outcome) and one point for correct rejections (correctly predicting a dirty outcome).

Results

The results of this experiment (Commons, Grotzer, & Davidson, in preparation) clearly show the effectiveness of reinforcement in contrast to nonreinforcing feedback on stage change. Here, mean sensitivity represents the formal-operational performance of detecting which causal relation holds in a particular problem and thereby predicts the outcome of a given episode. Sensitivity to casual relations is represented by non-normal d′ (Commons, Kantrowitz, Buhlman, Grotzer & Ellis, in press; Kantrowitz, Buhlman, & Commons, 1985). Mean sensitivity of 1.00 indicates the perfect detection of formal operational relationships, whereas a mean sensitivity of 0 indicates performance at the chance level, indicative of concrete operations. The mean sensitivity to causal relations across all subjects receiving feedback is plotted in the top panel of Fig. 7.1, with reinforcement (receiving valued points as well as feedback) in the bottom panel of Fig. 7.1. Without reinforcing points, even with verbal feedback as to correctness, there was no significant improvement across trials ($r^2 = .05$). With the reinforcing points, formal operational performance increased significantly ($r^2 = .59$, $p < .0005$).

Discussion

The laundry problem has shown that three conditions are necessary for stage transition. First, the subjects must be presented with a problem requiring performance at a stage higher than the subjects' present stage. The problem may be self-presented, but this delays the transition process tremendously. Second, the subjects must receive feedback from their attempts at solving the problem. Finally, the subjects must have their successful problem solutions reinforced.

The laundry problem also fulfills all the required conditions for intelligent pattern recognition. The environment in which the student is placed involves elements of regularity (a fixed causal relation), complexity (three distractor variables), and dynamism (change in variable values). The game is interactive with success or failure reported immediately so that the student takes part in an almost ideal feedback loop. Students' biological preferences have led to a complex chain of learned behavior. As a result of the learned behavior, the biological preference became manifested as psychological preferences for receiving a reinforcing point. Without the strong preference for the reinforcing point, no recognition of the causal pattern was acquired. The use of memory is clearly required

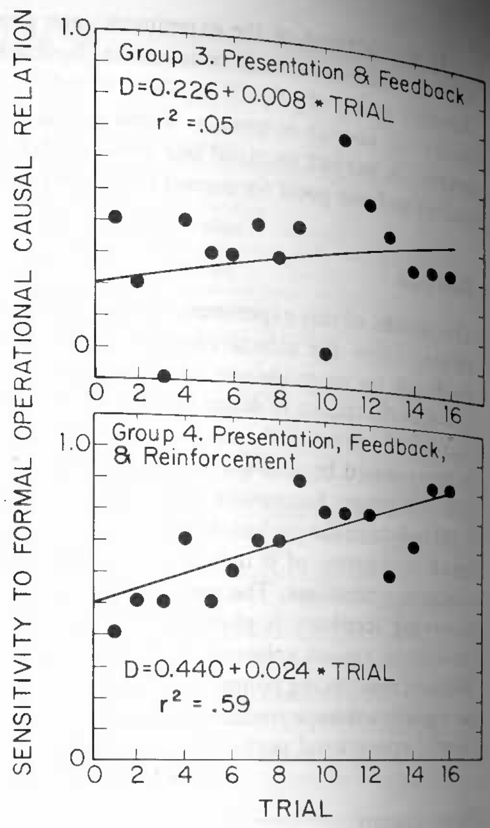

FIG. 7.1. Mean non-normal d', representing sensitivity to a causal relation versus trial. Each trial consisted of 10 predictions. In the top panel, "Correct" followed correct predictions, and "Incorrect," wrong ones. In the bottom panel, points as well as verbal feedback followed correct predictions.

because the students must play for more than one trial. Finally, that the students adapt and do so in a way that generates a new concept is shown next.

THE LAUNDRY PROBLEM AND LEARNING THEORY

Exploring the implications of these findings for AI research requires an understanding of the strategies used by the subjects in the laundry problem. Although the subjects may engage in a variety of activities while working on the laundry problem, here the following are examined: the information subjects used and some ways in which they used it; and the necessity for repeated opportunities to solve the problem, with feedback and reinforcement. To understand the development of strategies without appealing to subject's explanations, choice data were used (Kuhn & Brannock, 1977). Stein and Commons (1987; in preparation) show that verbal explanations and descriptions develop later than the nonverbal predictions.

TABLE 7.3
Three Overall Strategies for Making Predictions

Subjects Could Form Three Strategies to Make Predictions:
(1) pure guess (includes any manner of random choosing)
(2) object dependence (based on information from outside the experimental situation)
(3) direct matching (eight steps in transition from concrete to formal)

During the stage-change process, three types of strategies are observed. These types are outlined in Table 7.3. The concrete operational pattern of action consists of directly matching information in episodes to outcomes. This last strategy leads to formal operational matching, with subjects assessing which washing agents match outcomes in both the Test Episode and Informational Episode.

Piaget's Four Step Model of Stage Change

Because only one stage is under consideration here, Piaget's four-step model of stage change is most applicable to show how the direct matching pattern of behavior develops. Piaget's four-step, probabilistic model (Flavell, 1963) of equilibration illustrates the stage-change process. This model presents the steps through which a "new equilibrium" comes about. The four steps that are said to establish an equilibrium state are as follows: first, subjects attend to only one aspect of a problem; second, subjects begin to alternate between attending to one

TABLE 7.4
Piaget's Model of Equilibration as Seen in the Laundry Problem

A.	Subjects attend to washing agents in the test episode and try to find an Informational Episode with the most similar washing agents. The problem allows for a match of two or three agents. Subjects predict the outcome from that Informational Episode's set of washing agents. (Concrete)
B.	Subjects switch which washing agents they attend to in the Test Episode. They then find an Informational Episode with the most similar washing agents. Subjects predict the outcome from that Informational Episode set of washing agents.
AB.	Subjects attend to washing agents in the Test Episode and try to find two or more informational episodes with the most similar washing agents and the same outcome. The problem requires that they ignore at least two variables to find such a match. They predict outcome from these two or more episodes.
A with B.	Subjects first attend to the Informational Episodes. They look for a washing agent that always predicts clean. They then try to find that single washing agent in the Test Episode that matches. Although they use only on washing agent, it sometimes produces both clean and dirty clothes. Most often there is an increase in errors at this point for a short period of time. (Abstract)

not unintelligent encodings, take place (Campbell & Bickhard, 1986). For example, in isolating a variable in a situation like the laundry problem, it is a random element that will first bring the subject to attend to one variable alone. This requisite randomness is a learning factor that AI theory has largely neglected.

Some steps in the detection process are environmental; others are internal, in that they come about without any change in the external environment. The logic used to create representations is a property of those internal processes (Campbell & Bickhard, 1986).

The advantage of reducing the discrimination steps to these steps in respondent conditioning is that respondent conditioning is easily modeled in Artificial Intelligence. An examination of the complex type of conditioning seen in the laundry problem detailed previously will then clarify both the limitations of contemporary AI theory and suggest features of the learning process that AI might successfully exploit.

Classical and Operant Conditioning in the Discrimination Process

The detection steps outlined previously express a learning process known as operant conditioning. Operant conditioning can account for the success of reinforcement in inducing formal operational strategies as seen in the laundry problem. Operant conditioning involves providing reinforcement after subjects perform an act in a certain context. This reinforcement increases the likelihood that the subjects will perform the same act the next time they encounter an equivalent situation.

Commons and Armstrong-Roche (1984; also see Pear & Eldridge, 1984) argue that operant conditioning results from two classical conditioning steps in rapid succession. Several studies (Bindra, 1976, 1980; Gormezano, & Kehoe, 1982 for a review; Libet, 1985; see also Pear & Eldridge, 1984) provide support, albeit indirect, for this claim.

Classical (respondent) conditioning is illustrated by the following example. Suppose that getting a prize (a stimulus) always makes a certain subject excited (a response). Now suppose that an experimenter gives this subject a point leading to a prize (a second stimulus). If this sequence is repeated, the subject eventually will become excited as soon as the point is delivered. In effect, the point replaces the prize as reinforcer by its repeated association with the prize. Most respondent conditioning involves unconditioned responses that become "automatic" in just this way; that is, these responses occur without conscious thought. Operant conditioning takes place when the initial stimulus–response pairing of classical conditioning is followed by a second pairing that reinforces behavior elicited by the first.

In respondent conditioning, a stimulus must be predictive of a salient second stimulus (Commons, Herrnstein, & Wagner, 1983; Rescorla & Wagner, 1972)

for learning to take place.[11] This connection between the stimuli is also a requirement for the two respondent conditioning steps of operant conditioning. Learning will not proceed until the subject treats the critical stimuli in the conditioning situation as salient.

The laundry problem depicts this operant-conditioning process. As pointed out earlier, subjects were presented with the Informational Episodes and the first Test Episode. Assume that the following two outcomes have the following properties before the operant conditioning begins:.Hearing "correct" and receiving a point for their team together is an exciting reinforcer. These events have a history of leading to prizes. Hearing "incorrect" and receiving no point is a weak punisher. For operant conditioning to take place in this case, the decision preceding choice would have to be salient,[12] as would the relationship between the controlling ingredient and the state of the cloth. The Appendix presents a diagrammatic illustration of the process.

In the first conditioning step in the reduction, the "what to do" step, subjects come to attend to their internal decision.[13] The what-to-do step makes the cause of the operant response salient. That internal decision to make a given choice serves as an internal stimulus that elicits their choice. In the example, this choice is a correct or incorrect prediction of clean or dirty. On each occasion, they attend to this newly salient decision because it is predictive of a highly salient point reinforcer.[14] Thus the subjects learn what it is that they plan to do to obtain the reinforcer. This new decision to act in a certain way must compete with old decisions. In this instance, all parts of the new decision must compete with old preconceived notions about how to do laundry.

In the second conditioning step, the "when and where to do it" step, subjects discriminate among various types of external stimuli and distinguish the relevant stimuli. In this case, the stimulus to be discriminated is the relationship between the controlling ingredient and the state of the cloth. This relationship can be seen by examining the bottles of ingredients on the table and the state of the cloth placed next to them. The decision to predict clean or dirty, being highly salient, can be conditioned to correspond to the external stimuli that are already highly salient. This conditioning occurs because the relationship between the controlling ingredient and the state of the cloth is a stimulus that is predictive of a highly

[11]A stimulus is known to be salient if it regularly elicits some external response.

[12]The decision preceding choice and the events associated with that decision become salient to the subject when it regularly elicits an external response.

[13]This step is relatively new in conditioning theory. For examples see Bindra (1976) and Commons and Armstrong-Roche (1984). The decision is internal because, without electrophysiological measurement, it is not directly observable whereas the choice that follows the internal decision is directly observable. The decision is the transient plan that precedes the choice.

[14]Some of the salient characteristics elicited by "correct" and "point" come to be elicited by the internal decision.

salient stimulus—the decision to choose clean or dirty. If this decision were not salient, this second step would not proceed. Likewise, if each ingredient were not treated as a salient stimuli to predicted cloth outcomes, this second step would not proceed. The conjunction of these salient stimuli, the decision and the relationship between the control washing agent and the outcome of the cloth, has been seen in the laundry problem to be highly successful in promoting formal operational pattern recognition, as seen earlier in Fig. 7.2.

PATTERN RECOGNITION AS A MODEL FOR ARTIFICIAL INTELLIGENCE

Six requirements for intelligent pattern recognition were listed in Table 7.1 as the basis for both a critique of current AI theory and a model for more successful attempts at creating intelligent machines. Having explored in an example the process of pattern recognition in the human sphere, and the mechanisms that promote it, it is now possible to describe against this background the failure of AI to model appropriate systems and provide adaptive mechanisms.

Current AI theory often attempts to model systems comprised of programs in specific-task domains. There are two reasons for this. First, it seems too difficult to construct models accounting for all intelligent behavior; it seems more productive to construct systems that can solve tasks in a limited domain and then generalize the systems to include wider ranges of behavior. Second, the lesson AI drew from the failure of learning-based systems in the 1950s and 1960s was that programs had to have already extensive knowledge bases before they could acquire more knowledge in a way commensurate with human learning (Michalski, Carbonell, & Mitchell, 1983). Both these viewpoints have problems.

First, as already argued, not all task domains are capable of supporting intelligent behavior. The laundry problem forms a domain in which intelligent pattern recognition could occur, but only because it fulfilled the six conditions in Table 7.1. Compare this with DART (Genesereth, 1984), a device-independent diagnostician that employs common AI problem-solving techniques (Davis, 1984; de Kleer, 1984; Winston, 1984). DART accepts as input a complete model of the device structure (parts and interconnections) and expected behavior (relations between device input, internal state, and output). Then when provided with an illegal input–output pair, DART uses a theorem-proving technique (resolution residue) to generate a suspected broken component, generate an input–output pair to test that component, and then perform the test.

Although successful as a diagnostician, DART does nothing more than search a solution space for one that works. Two questions must be answered. First, can this behavior be generalized to other domains? Second, is this intelligent problem solving? The first can be answered positively, in much the same sense that one can build expert systems for different areas of expertise. The second cannot, because novelty within the chosen domain is disallowed. For example, what happens if the device model provided by the programmer is wrong? The hard problems are defined to lie outside the domain, effectively leaving the system in a state of stasis. Other AI successes, like XCON (McDermott, 1982), DENDRAL (Lindsay, Buchanan, Feigenbaum, & Lederberg, 1980), and MYCIN (Shortliffe, 1976), share the same fault.

In fact, many of the common problem-solving paradigms in AI are not conducive to handling novelty. Goal-reduction is a method of recursively decomposing a problem into subgoals until a subgoal can be equated with some known procedure or knowledge. Means–end analysis, a control paradigm, finds a solution path through a graph by moving to the node or subproblem that most reduces the distance between the program's current state and goal state. Constraint-propagation reduces labels of nodes on a graph describing the solution space. The space of tasks that can be solved by these techniques and other common AI techniques (see Winston, 1984) is exactly the space of tasks that can be solved by the programmer-supplied primitives. In essence, these techniques are all methods for searching a knowledge base, and, as shown earlier, search does not suffice. A program can be given knowledge of the set of integers and arithmetic operations, but unless it has a mechanism for generalization—learning new concepts—it will not be able to solve an algebraic equation containing a variable, because the notion of variable must be learned and is not reducible to notions of specific integers.

The second argument AI theory maintains against emphasizing learning is that previous systems failed to acquire knowledge because learning requires some previous knowledge. This is in effect a belief in creationism, for it claims that intelligent mechanisms cannot bootstrap themselves from ground zero, in direct contrast with the evidence of evolutionary theory. While humans do begin life with an impressive array of hardwired cognitive abilities, humans also have to have many years of experience before their cognitive abilities approach maturity.

To take another perspective, consider the mapping M that takes machine input to machine output. Then M can be represented by the composition of a series of submappings M1, M2, M3, . . . where M1 might be a low-level vision processor (e.g., the lateral-geniculate nucleus), M2 might be a higher level vision processor (e.g., area 17), and M3 might be a recognition function, and so on. Each of these submappings must take as input the output of the previous mapping. One therefore cannot arbitrarily choose representations for the knowledge processed by intelligent systems. The representations must have a place in the system as a whole. Learning to manipulate arbitrarily chosen knowledge bases for arbitrarily chosen tasks is by definition independent of the submappings comprising M. In addition, there is no reason to suppose that knowledge representations chosen are correct.

In fact, to truly adapt a system to a particular environment, it must not start with any preconceptions about the environment. Recent research indicates that there may exist natural modes, or natural categories (Commons, Herrnstein, & Wagner, 1983; Herrnstein & Loveland, 1964; Herrnstein, Loveland, & Cable, 1976) that reflect a good tradeoff between the ease of categorizing and the amount of information gained by placing an object in the category. These modes are environment specific.

A major difference between animals and machines is that animals derive representations of information from experience, whereas present AI machines simply have been encoded by their programmers. A machine that works with encodings can only generate copies or translations (Fodor, 1975, 1981), which tend to preserve information and do more of the same task. Doing more of the same clearly involves no learning (Campbell & Bickhard, 1987).

Yet another argument against beginning with much knowledge is that the amount of data impinging on humans, and by extension on any machine in a sufficiently complex environment, is so vast that a large portion of the information must be discarded. The machine must learn to attend to only the most salient aspects of the input. The knowledge base that the machine develops will therefore reflect the most salient information and that cannot be decided independently of the response of the environment to machine input.

Obeying the principle of adaptation poses perhaps the most difficult challenge to AI theory. As suggested previously, the degree of adaptibility demanded within a domain is an index of the domain's usefulness in a divide-and-conquer approach to intelligence. Too often, systems are designed to recognize pre-established regularity. Categorizing typewritten characters is one task domain that has been solved by statistical and feature-based methods, but it is significant that handwritten character recognition, an only slightly less trivial domain, has only been solved with learning techniques.

Marr's theory of competence (Marr, 1982), which has served as a framework for much of recent vision research, offers some insight into the AI perspective. This framework exploits Chomsky's distinction between performance and competence in its use of a three-tiered hierarchy to relate the different possible levels of explanation. The problem of vision may be taken as solved when it can be completely described at the level of: (1) a theory of competence describing the goals, assumptions, inputs, and transforms necessary to achieve the goals; (2) representations of the input, output, and intermediate data, and algorithms instantiating the necessary transforms; and (3) implementation, i.e., neural or electronic circuitry instantiating the algorithms and representations. Together the levels of description dictate performance, which can be matched against the results of psychophysics and neurophysiology to assess the correlation between theory and human behavior. If the failures and the successes of the theory match human failures and successes, then the correlation is taken to be high. This is the essential idea of the Turing Test, though strengthened by the applications of natural constraints.

Whereas this approach engenders goal-directed vision systems, the very goal ascribed to the system by the theory of competence must correspond to the goal of the human system, and it is clear that the only goal from which system behavior is generated is by definition. Though machine preferences are constant over time, environmental conditions are not, and so the mechanisms that are needed to satisfy those demands must change. Theories of competence that suppose a fixed goal fail adaptability criteria. In the particular case of low-level vision, Marr's level of explanations may hold, but only because low-level vision consists of hardwired primitives that cannot adapt. Theories of competence work for low-level hardwired processes, but not for higher processes.

Complementing the process of self-organization (adaptation) is the process of self-stabilizing (Vaughan & Herrnstein, 1987). Self-monitoring mechanisms must be developed that compile or hardwire frequently useful behaviors. The gain in response time increases the likelihood of success in a highly dynamic environment. Additionally, hardwired or compiled structures are the ones that would be most useful to transfer to copies of the machine that are to exist in the same environment.

A major objection of AI to learning theory is that machines that construct their own representations may not tell us anything about how humans solve their problems, or, more generally, do not explain how thought actually works. This objection rests on the assumption that there is some particular representation of knowledge or particular way of manipulating knowledge that underlies all intelligence. There is no clear reason to assume this, and many reasons not to. It may be true that there is a common adaptive mechanism and that the learning mechanism constrains the representations available. But in this case one has to tackle the problem of adaptation first, not second as AI theorists would like to do.

CONNECTIONIST MODELS

The recent connectionist approach to intelligence does address the problem of adaptability. Whereas previous perceptron models could only compute linear transformations of input data, current neural nets can construct internal representations of input data that are not linearly related to the input. Rumelhart (1987) has developed a distributed-knowledge, supervised learning technique called back propagation that is successful at a wide range of tasks, such as reading English out loud and discriminating submarine echoes from noise in sonar data (Sjenowski, 1986). Fukushima, Miyake, and Ito (1983) have developed a localized-knowledge, supervised network that displays deformation and position invariance in learning to categorize handwritten arabic numerals. Other variants abound (Anderson, 1983; Barto, Sutton, & Anderson, 1983; Grossberg, & Mingolla, 1985a,b; Hopfield, 1982, 1984). These techniques, however, in general do not yet exploit the operant-conditioning mechanism that is so successful in promoting stage change. One exception is Carpenter and Grossberg (this volume), who have developed a network that does employ operant conditioning.

CONCLUSION

This analysis has linked intelligence to pattern recognition in direct contrast to AI theorists' linking intelligence to knowledge representation and manipulation. Knowledge representation should arise out of experience and not be impressed on the system by the outside assistance of the programmer. Knowledge manipulation is the result of operant-conditioning steps that map perception to action. Though it is unfair to conclude that AI theory has been unproductive to date, it is fair to argue that AI's progress toward the goal of intelligent behavior will be impeded as long as AI assumes the perspective that McCarthy first put forward (cited in Minsky, 1968): "We base ourselves on the idea that in order for a program to be capable of learning something it must first be capable of being told it." As shown here, McCarthy's claim is exactly wrong. Acquisition and exploitation of knowledge are not independent but complexly interwoven.

REFERENCES

Anderson, J. A. (1983). Cognitive and psychological computation with neural models. *IEEE, SMC, 13* (5), 799–814.

Barto, A. G., Sutton, R. S., & Anderson, C. W. (1983). Neuron-like adaptive elements that can solve different learning control problems. *IEEE, SMC, 13* (5), 834–846.

Bell, C. C. (1981). An efference copy which is modified by reafferent input. *Science, 214*, 450–452.

Bindra, D. (1976). *A theory of intelligent behavior*. New York: Wiley.

Bindra, D. (1980). *The brain's mind: A neuroscience perspective on the mind–body problem*. New York: Gardner Press.

Campbell, R. L., & Bickhard, M. H. (1986). Knowing levels and developmental stages. *Contributions to human development*.

Campbell, R. L., & Bickhard, M. H. (1987). A deconstruction of Fodor's anticonstructivism. *Human Development 30*(1), 48–59.

Case, R. (1985). *Intellectual development: Birth to adulthood*. Orlando: Academic Press.

Chase, S., & Heinemann, E. G. (in press). Recognition, classification and identification of signals by people and pigeons: Effects of limited memory. In M. L. Commons, M. C. Davison, J. A. Nevin, (Eds.), *Models of behavior, 10, Signal detection*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Commons, M. L., & Armstrong-Roche, M. (1984, August 24–28). *Some respondent processes underlying operant conditioning*. Paper presented at the 92nd Annual Convention of the American Psychological Association, Toronto.

Commons, M. L., Frome, K. W., & Armstrong-Roche, M. (in preparation). *The sense of free will*.

Commons, M. L., Kantrowitz, S., Buhlman, R. A., Ellis, J. B., & Grotzer, T. A. (in press). Immediately detecting a causal-relation stimulus embedded in similarly appearing possible causal relations. In M. L. Commons, M. C. Davison, J. A. Nevin, (Eds.), *Models of behavior, 10, Signal detection*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Commons, M. L., Grotzer, T. A., & Davidson, M. N. (in preparation). *The necessity of reinforcing problem solutions for transition to formal operations: An examination of Piaget's equilibration theory of stage change*.

Commons, M. L., Handel, A. N., & Richard, D. C. & Grotzer, T. A. (in preparation). *The development of strategies during the transition to formal operations*.

Commons, M. L., Herrnstein, R. J., & Wagner, A. R. (Eds.). (1983). *Quantitative analyses of behavior: Vol. 4. Discrimination processes*. Cambridge, MA: Ballinger.

Commons, M. L., Miller, P. M., & Kuhn, D. (1982). The relation between formal operational reasoning and academic course selection and performance among college freshmen and sophomores. *Journal of Applied Developmental Psychology, 3*, pp. 1–10.

Commons, M. L., & Richards, F. A. (1984a). A general model of stage theory. In M. L. Commons, F. A. Richards, & C. Armon (Eds.), *Beyond formal operations: Vol. 1. Late adolescent and adult cognitive development* (pp. 120–140). New York: Praeger.

Commons, M. L., & Richards, F. A. (1984b). Applying the general stage model. In M. L. Commons, F. A. Richards, & C. Armon (Eds.), *Beyond formal operations: Vol. 1. Late adolescent and adult cognitive development* (pp. 141–157). New York: Praeger.

Commons, M. L., Richards, F. A., & Armon, C. (Eds.). (1984). *Beyond formal operations: Vol. 1. Late adolescent and adult cognitive development*. New York: Praeger.

Davidson, M. N. (1983). *The effect of repeated presentation of a problem on the transition from concrete to formal operations: A test of Piaget's theory of stage change*. Unpublished B.A. honors thesis, Harvard University.

Davidson, M. N., & Commons, M. L. (1983, April 6–9). *How repeatedly presenting a formal operational problem with no feedback affects the transition from late concrete through beginning formal to formal operations: An examination of Piaget's equilibrium theory of stage change*. Paper presented at the annual meeting of the Eastern Psychological Association at Philadelphia.

Davis, R. (1984). Diagnostic reasoning based on structure and behavior. In D. G. Bobrow (Ed.), *Qualitative reasoning about physical systems*. Cambridge, MA: MIT Press.

de Kleer, J. (1984). How circuits work. In D. G. Bobrow (Ed.), *Qualitative reasoning about physical systems*. Cambridge, MA: MIT Press.

Dore, F. Y., & Dumas, C. (1987). Psychology of animal cognition: Piagetian studies. *Psychological Bulletin, 102*, 219–233.

Fischer, K. W. (1980). A theory of cognitive development: The control and construction of hierarchies of skills. *Psychological Review, 87*, 477–531.

Flavell, J. H. (1963). *The developmental psychology of Jean Piaget*. New York: Van Norstrand.

Fodor, F. (1975). *The language of thought*. New York: Crowell.

Fodor, F. (1981). The present status of the innateness controversy. In J. Fodor (Ed.), *Representations*. Cambridge, MA: MIT Press.

Fukushima, F., Miyake, S., & Ito, T. (1983). Neocognitron: A neural network model for a mechanism of visual perception. *IEEE, SMC, 13* (5), 826–834.

Genesereth, M. R. (1984). The use of design descriptions in automated diagnosis. In D. G. Bobrow (Ed.), *Qualitative reasoning about physical systems*. Cambridge, MA: MIT Press.

Gormezano, I., & Kehoe, E. J. (1982). Associative transfer in classical conditioning to serial compounds. In M. L. Commons, R. J. Herrnstein, & A. R. Wagner (Eds.), *Quantitative analyses of behavior: Vol. 3. Acquisition* (pp. 297–322). Cambridge, MA: Ballinger.

Grossberg, S., & Mingolla, E. (1985a). Neural dynamics of form perception: Boundary completion, illusory figures, and neon color spreading. *Psychological Review, 92*(2), 173–211.

Grossberg, S., & Mingolla, E. (1985b). Neural dynamics of perceptual grouping: Textures, boundaries, and emergent segmentations. *Perception & Psychophysics, 38*(2), 141–171.

Hebb, D. O. (1949). *The organization of the brain*. New York: Wiley.

Herrnstein, R. J. (1970). On the law of effect. *Journal of the Experimental Analysis of Behavior, 13*, 243–266.

Herrnstein, R. J. (1982). Melioration as behavioral dynamism. In M. L. Commons, R. J. Herrnstein, & H. Rachlin (Eds.), *Quantitative analyses of behavior: Vol. 2. Matching and maximizing accounts* (pp. 433–458). Cambridge, MA: Ballinger.

Herrnstein, R. J. (in press). *Levels of categorization.*

Herrnstein, R. J., & Loveland, D. H. (1964). Complex visual concept in the pigeon. *Science 146*, 549–551.

Herrnstein, R. J., Loveland, D. H., & Cable, C. (1976). Natural concepts in pigeons, *Journal of Experimental Psychology: Animal Behavior Process, 2*, 285–311.

Herrnstein, R. J., & Vaughan, W., Jr. (1980). Melioration and behavioral allocation. In J. E. R. Staddon (Ed.), *Limits to action: The allocation of individual behavior* (pp. 143–176). New York: Academic Press.

Hobbes, T. (1650). *Humaine nature: Or the fundamental elements of policie.* London.

Hobbes, T. (1651). *Leviathan, or the matter, forme and power of a commonwealth, ecclesiasticall and civill.* London.

Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational properties. *Proceedings of the National Academy of Sciences, 79*, 2554–2558.

Hopfield, J. J. (1984). Neurons with graded response have collective computational properties like those of two-state neurons. *Proceedings of the National Academy of Sciences, 81*, 3088–3092.

Inhelder, B., & Piaget, J. (1958). *The growth of logical thinking from childhood to adolescence.* New York: Basic Books.

Jennings, H. S. (1906, 1976). *Behavior of the lower organisms.* Bloomington: Indiana University Press.

Kantrowitz, S., Buhlman, R. A., & Commons, M. L. (1985, April). *Application of signal-detection theory to the measurement of stage of development found in performance on formal-operational isolation-of-variables problems.* Paper presented at the annual meeting of the Eastern Psychological Association at Boston.

Kehoe, E. J. (in press). A layered network model of associative learning: Learning-to-learn and configuration. *Psychological Review.*

Kuhn, D., & Brannock, J. (1977). Development of the isolation of variables scheme in experimental and "natural experiment" contexts. *Developmental Psychology, 13*, (1), 9–14.

Libet, B. (1985). Unconscious initiative and the role of the conscious will in voluntary action. *The Behavioral and Brain Sciences, 8*, 529–539.

Lindsay, R., Buchanan, B. G., Feigenbaum, E. A., & Lederberg, J. (1980). *Applications of artificial intelligence for chemical inference: The DENDRAL project.* New York: McGraw-Hill.

Linn, M. C., Chen, B., & Thier, H. D. (1976). Personalization in science: Preliminary investigation at the middle school level. *Instructional Science, 5*, 227–252.

Linn, M. C., Chen, B., & Thier, H. D. (1977). Teaching children to control variables: Investigation of a free choice environment. *Journal of Research in Science Teaching, 14* (3), 249–255.

Marr, D. (1969). A theory of cerebellar cortex. *Journal of Physiology, 202*, 437–470.

Marr, D. (1982). *Vision.* San Francisco: W. H. Freeman.

Matsubara, J. A. (1981). Neural correlates of a nonjammable electrolocation system. *Science, 211*, 722–725.

McClelland, J. L., Rumelhart, D. E., & the PDP Research Group. (1985). *Parallel distributed processing: Explorations in the microstructure of cognition. Volume 2: Psychological and biological models.* Cambridge, MA: MIT Press.

McDermott, J. (1982). R1: A Rule-based configurer of computer systems. *Artificial Intelligence, 19*(1).

Michalski, R. S., Carbonell, J. G., & Mitchell, T. M. Eds. (1983). *Machine learning: An artificial intelligence approach.* Los Altos, California: Morgan Kaufmann.

Minsky, M. (1968). *Semantic information processing.* Cambridge, MA: MIT Press.

Newman, E. A., & Hartline, P. H. (1981). Integration of visual and infrared information in bimodal neurons of the rattlesnake optic tectum. *Science, 213*, 789–791.

Pascual-Leone, J. (1980). Constructive problems for constructive theories: The current relevance of Piaget's work and a critique of information-processing stimulation psychology. In R. Kluwe & H. Spada (Eds.), *Developmental models of thinking.* New York: Academic Press.

Pate, R. H. Jr. (1970). Student Expectations and later expectations of a university enrollment. *Journal of College Student Personnel, 11*, 458–462.

Patrick, E. A. (1972). *Fundamentals of pattern recognition.* Englewood Cliffs, NJ: Prentice-Hall.

Pear, J. J., & Eldridge, G. D. (1984). The operant-respondent distinction: Future directions. *Journal of the Experimental Analysis of Behavior, 42*, 453–467.

Peitgen, H. O., & Richter, P. H. (1986). *The beauty of fractals.* New York: Springer–Verlag.

Piaget, J. (1952). *The origins of intelligence in children.* New York: International Universities Press.

Pollen, R. A., & Ronner, S. F. (1983). Visual cortical neurons as localized spatial frequency filters. *IEEE Transactions on Systems, Man, and Cybernetics, 13*(5), 907–915.

Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and non-reinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning* (Vol. 2). New York: Appleton-Century-Crofts.

Resnikoff, H. L. (1985). *The illusion of reality: Topics in information science.* New York: Springer-Verlag.

Rumelhart, D. E., McClelland, J. L., & the PDP Research Group. (1985). *Parallel distributed processing: Explorations in the microstructure of cognition. Volume 1: Foundations.* Cambridge, MA: MIT Press.

Shortliffe, E. H. (1976). *MYCIN: Computer-based medical consultations.* New York: Elsevier.

Siegler, R. S., Liebert, D. E., & Liebert, R. M. (1973). Inhelder and Piaget's pendulum problem: Teaching preadolescents to act as scientists. *Developmental Psychology. 9*(1), 97–101.

Sjenowski, T. (1986, October). *Paper presented at Neural Networks and Neuromorphic Systems conference*, Boston.

Skinner, B. F. (1969). *Contingencies of reinforcement.* New York: Appleton-Century Crofts.

Spencer, H. (1870). *Principles of psychology* (2nd ed.). London: Longman.

Stein, S. A., & Commons, M. L. (1987, May). *Verbal description and explanation appear lower in stage than nonverbal performance on a formal operational task.* Paper presented at the meeting of the Jean Piaget Society, Philadelphia.

Stein, S. A., & Commons, M. L. (in preparation). *Verbal description and explanation appear lower in stage than nonverbal performance on a formal operational task.*

Thorndike, E. L. (1898). Animal intelligence: An experimental study of the associative process in animals. *Psychological Review Monographs, Supplement 2*(8), 28–31.

Thorndike, E. L. (1911). *Animal intelligence.* New York: Macmillan.

Vaughan, W., Jr., & Herrnstein, R. J. (1987). Stability, melioration, and natural selection. In L. Green & J. H. Kagel (Eds.), *Advances in behavioral economics* (Vol. 1, pp. 185–215). Norwood, NJ: Ablex.

Winston, P. H. (1984). *Artificial intelligence* (2nd ed.). New York: Addison-Wesley.

APPENDIX

A Formalization of the Reduction: Initial Steps

The stimulus–response pairings of classical conditioning can be used to reduce the more complex, two-stage process of operant conditioning to a single ex-

TABLE 7.6
Symbol Definitions Without Their Subscripts for the Reduction

(a) cr = an internal response for which there is no observable causal sequence
(b) CR = conditioned response
(c) CS = conditioned environmental stimulus that has come to elicit a response
(d) p( ) = probability of
(e) s = stimulus in the environment to be conditioned
(f) S^D = discriminative stimulus
(g) t = time
(h) us = internal stimulus that precedes an operant response. This is the short form of us_1
(i) US = environmental stimulus that already elicits a response
(j) ... = first event has a second event programmed to follow
(k) - = one event elicits another
(l) ---> = one event elicits another
(m) ===> = leads to, over a number of presentations
(n) |--o--| = vertical dashed lines identify the specific elements to be paired. The line points to the single letter which will be paired.
E_1 E_2
(o) o o o -----o----- = pairing operations
(p) o = pairing operation (i.e., the existence of one stimulus is dependent on the occurrence of the other stimulus)
(q) p(|) = conditional probability, probability of one event given that another has occurred
(r) uppercase italicized letters = presently observable events
(s) lowercase italicized letters = internal events that are potentially observable
(t) ≤ = less than or equal to
(u) ≥ = greater than or equal to
(v) , = and
(w) X_1/X_2 = an event with properties X_1 and X_2

pression. The detection process outlined in Table 7.5 describes the same process in linear form. Step 2 in Table 7.5 corresponds with pairing #1a in expression 3; steps 3 and 4 correspond with pairing #2. The symbols used to diagram this reduction are defined in Table 7.6.

Diagram of the Proposed Reduction

Respondent conditioning is written in traditional form as follows:

$$S \quad o \quad US \dashrightarrow UR \qquad (1)$$

After conditioning:

$$CS \dashrightarrow CR \qquad (2)$$

To illustrate: *S* could be the experimental situation in step 1 and *US* could be the reinforcers and punishers. The situation comes to elicit the same transformations of the responses to the punishers and reinforcers.

Thus in respondent conditioning only one pairing takes place, relating the *US* (reinforcer) to the *S*/*CS*. On the other hand, in operant conditioning two other critical pairings occur. Thus there are three pairings altogether: the two additional critical pairings, and the pairing of the US_2 with the neutral S_1/S^D as shown in Expression 3, in terms of the laundry problem.

pairing #1a
[*cr-us*] → *R*. . . . *Q* . . . . . US_2/S^{R+}
internal events (decision) — saying clean when clean — point & "correct"
pairing # 2 *o* . . . *o* pairing #1b
S_1/S^D
critical washing agent (3)

These three pairings result in the following chain of events:

$$CS/S^D - (us - CR_2) - (R_1) \ldots\ldots US_2/S^{R+}$$

critical washing agent — salient internal event — saying clean when clean — point & "correct" (4)

After pairing #1a (step 2) and pairing #2 (step 3 and 4), the introduction of the critical washing agents CS/S^D elicits the response of excitement $CR_{2(R)}$, which has previously been elicited as part of the little *us* complex, (us-CR_2/CS_2). The probability of these responses $CR_2(R)$ has increased because, in pairing #2, the complex that elicits the response has been paired with the environmental stimulus S_1. These responses CR_2 were originally elicited only by the delivering of the prize (US_2) and *R* was originally elicited only by *us*.