

Second Edition

ENCYCLOPEDIA OF ARTIFICIAL INTELLIGENCE

**Volume 1
A-L**

Awarded
American Library Association's
Outstanding Reference Source
Association of American Publishers Award
Best New Professional and Scholarly Publication

Stuart C. Shapiro
Editor-in-Chief

Machine Learning

Jeffrey C. Schlimmer and Pat Langley,
LEARNING, MACHINE, 785-805

From Shapiro, Stuart C., Editor-in-Chief,
Encyclopedia of Artificial Intelligence, 2nd
Ed., John Wiley & Sons, Inc., NY, 1992.
"Copyright 1992 by John Wiley & Sons, Inc.
This material is reproduced with permission
of John Wiley & Sons, Inc."

LEARNING, MACHINE

One central insight of artificial intelligence is that expert performance requires domain-specific knowledge. Machine learning is the subfield of AI that studies the automated acquisition of such knowledge. The aim is intelligent systems that learn—that is, that improve their performance as the result of experience (see EXPERT SYSTEMS). Researchers study learning for a variety of reasons: to discover general principles of intelligence, to better understand human learning, and to automate the process of knowledge acquisition. However, the discipline is united in its concern with mechanisms for learning, and methods proposed with one goal in mind often serve other purposes equally well.

Despite its separate identity, there are two important senses in which machine learning is an integral part of the larger AI field. First, learning researchers must consider central AI issues of knowledge representation (qv), memory organization, and performance. Second, learning can occur in any domain requiring intelligence, whether the basic task involves classification, problem solving, reasoning, natural language processing, or vision. Thus, one can view machine learning as more a framework for AI research and development than as a subfield of AI.

This article reviews the current state of machine learning. It begins by identifying relevant dimensions along which learning systems differ. Next, it describes five of the different paradigms for machine learning that have emerged over the past two decades:

- Connectionist (neural network) learning methods
- Genetic algorithms and classifier systems
- Empirical methods for inducing rules and decision trees
- Analytic learning methods
- Case-based approaches to learning

Despite differences in the representations and algorithms used by these approaches, all aim to improve performance through the acquisition of knowledge from experience. Moreover, they must address many similar issues, and they share many open problems. After examining each of these paradigms, this article will consider some recent developments in the field as a whole and their potential impact on future research and applications.

One related area that will not be examined is knowledge acquisition (qv) (eg, Marcus, 1988). Like machine learning, this subfield of AI is concerned with improving performance through the storage of domain-specific knowledge. However, research on knowledge acquisition focuses primarily on direct interaction with an expert by asking questions, taking advice, and the like. In contrast, most work on machine learning emphasizes the acquisition of knowledge through experience with an external environment or experience with internal problem-solving traces.

DIMENSIONS OF MACHINE-LEARNING RESEARCH

Before delving into recent research on machine learning, let us identify some dimensions along which alternative learning paradigms vary. Six issues are discussed below: the representation of experience, the representation of acquired knowledge, the performance task, the difference between supervised and unsupervised learning, the distinction between incremental and nonincremental learning, and the complementary notions of induction and explanation.

Representation of Experience

Most learning is based on experience, and this requires a representation for the experiential input given to the learning system. Researchers have employed three broad classes of data structures, each with different representational power. The simplest describes each individual experience as a list of binary *features*, each corresponding to the presence or absence of some aspect of the environment. For example, a particular symptom (eg, jaundice) may be present or absent for a given patient. Connectionist and genetic algorithms typically employ feature-based input.

A second scheme assumes a known set of *attributes*, each having a set of mutually exclusive values. Thus, one

might describe an object as blue or red, but not both at the same time. Some attributes take on numeric values (eg, length, weight). Attribute-value representations are typically used by empirical methods, such as those for rule induction and decision-tree construction.

A final approach employs relational or structural representations. These describe relations between two or more objects, such as the fact that object *A* is above object *B* (ie, arbitrary predicates). Such structural information can be stated in many formalisms, including predicate logic (see LOGIC, PREDICATE) and semantic networks (qv). Although relational schemes have significantly more expressive power than the other representations, they also introduce significant complexity into the matching process, and this affects learning methods that use them. Research on analytic learning typically deals with relational data structures, as does some work on empirical rule learning and case-based reasoning.

Representation of Acquired Knowledge

All learning systems acquire new knowledge, and they must represent that knowledge in some fashion. The choice of knowledge representation is central in that it strongly influences choices of performance and learning components. In fact, the five learning paradigms to be discussed in the next main section are distinguished largely in terms of representational issues.

One choice is whether to store individual, concrete experiences or only abstractions based on these data. Most work on machine learning has taken the latter path and attempts to move beyond the data to build more general knowledge structures that summarize previous experience. Of course, hybrid approaches are possible; these are often called *case-based* approaches and are discussed in a latter section (see also REASONING, CASE-BASED). Note that methods that attempt to store only concrete experiences must at some point generalize beyond the data if they hope to apply to new experiences. Thus, the real issue is whether such abstraction occurs aggressively as new experiences are stored (as in most approaches) or lazily as needed during knowledge access (as with case-based approaches).

For abstract knowledge structures, another choice is whether to use a logical, discrete formalism or one that involves numeric, continuous information. This is primarily a choice between using logical connectives or numeric ones, respectively. Empirical rule learning and analytic learning methods have predominantly used the first path, whereas connectionist systems have relied on the second. Genetic algorithms can be profitably viewed as combining these two formalisms, as can recent work on probabilistic learning. A related choice is whether to use an aggregated, coarse-grained representation or a finer-grained one. These options are sometimes referred to as symbolic and subsymbolic representations, respectively. Conventional wisdom associates the former scheme with empirical and analytic methods, and the latter with connectionist and genetic approaches. However, as we will argue in the subsection on Relationships Among the Paradigms, these associations are more myth than fact.

The Performance Task

Learning involves improvement in performance, and thus learning cannot occur in the absence of some performance task. The vast majority of machine-learning research has focused on two broad classes of domains: classification and problem-solving (qv). Each has led to different concerns, distinct insights, and unique applications.

Intelligent agents are repeatedly confronted with the need to classify or label their experience. For example, on encountering certain symptoms, a doctor diagnoses a specific disease. The generic task can be easily stated: given some description of an experience, along with a set of known classes and their descriptions, assign that experience to one or more classes. The most frequent successes in expert systems have involved just such classification or diagnostic tasks (Davis and Lenat, 1982). Much of the work on machine learning has focused on learning descriptions for classification.

However, an intelligent agent must also be able to solve novel problems and formulate plans. For example, when going to a meeting, an agent devises a path from its current location to the target site. The generic task can be stated as: given some desired state or goal, find some sequence of actions that takes you from the current state to the desired one. Navigation and robot manipulation (see ROBOT MANIPULATORS) are the most obvious applications of problem-solving techniques, but there are many others, including scheduling. This area has been less popular within machine learning, but much of the recent work has focused on acquiring knowledge to improve the speed and quality of methods for planning and reasoning.

Finally, intelligent agents also participate in other high-level behaviors that may not fit neatly into either the classification or problem-solving paradigms. Such behaviors include the design of new artifacts, communication with other agents, and the control of motor effectors. To date, the machine learning community has placed little emphasis on these areas, but because of their potential, this will undoubtedly change in the future.

Supervised and Unsupervised Learning

Another dimension that influences learning is the degree of supervision. In some cases, a tutor or domain expert may be present to give the learner immediate feedback about the appropriateness of its behavior. This situation is typically called *supervised* learning. In other cases, an *unsupervised* learner may have to fend for itself. Here the learner has little or no external guidance in building knowledge structures or composing solutions to problems; at most, the environment provides coarse-grained feedback about the learner's overall effectiveness on the task at hand.

Early work in machine learning focused on the simpler supervised task, and it continues to receive considerable attention. There are good practical reasons for this interest. Although experts generally have poor ability to introspect about their domain knowledge, they are much better at providing examples of correct and incorrect behavior.

Both forms of learning can occur in many contexts, though they take on different forms in different domains.

In classification domains, the supervised task is usually called *learning from examples*, and the individual training experiences are called *instances*. This task can be stated as:

- Given: A set of instances (eg, patient symptoms) and their associated classes (eg, diseases).
- Find: A general description for each class that matches only its instances.

This basic task has been examined within all the major paradigms of machine learning and, until recently, the vast majority of research papers have dealt with this topic. However, there is now also considerable interest in unsupervised concept learning (qv), in which the learner must decide for itself not only the class in which it should place each experience, but also the number of such classes. This has been called clustering (qv) (Fisher and Langley, 1985).

In problem-solving domains, supervised learning occurs when a tutor is available to suggest the correct operator or subgoal at each point in the search for a solution. Methods that operate in this context are sometimes called *learning apprentices*, and they show promise as a way of extracting knowledge from an expert (Mitchell and co-workers, 1985) in problem-solving and design domains. However, more research in this area has focused on unsupervised learning, in which the agent must distinguish desirable actions from undesirable ones for itself. This has been called the credit assignment problem (Sleeman and co-workers, 1982), and researchers have explored a variety of responses across a number of paradigms.

Incremental and Nonincremental Learning

Some learning algorithms process experiences one at a time, whereas others process a large set of experiences at once. The former class is often called *incremental* and the latter, *nonincremental*. However, examining only the surface behavior of a system can be misleading, as an inherently nonincremental system can always be run in incremental mode and vice versa. For example, if one lets a nonincremental technique retain exact copies of previous experience in memory, it can process the first experience, then process the first two experiences, then process the first three experiences, and so forth. Similarly, one can take an incremental method and iteratively run it through the same set of experiences again and again. A better definition of incremental involves the number of previously seen experiences the method must reprocess during learning (Schlimmer and Fisher, 1986).

Both nonincremental and incremental approaches have advantages. The former can collect statistics about all the data, giving a learning method more information on which to base its decisions. On the other hand, incremental methods tend to be more efficient and, in principle, can deal with much larger data sets. Much of recent machine-learning research has focused on incremental approaches, though there are some notable exceptions in the literature on connectionist methods and on the induction of rules and decision trees.

Inductive and Analytic Learning

A sixth important dimension relates to the type of learning used to acquire knowledge. In one class of approaches, inductive learning methods formulate knowledge based mainly on observed data. Empirical, genetic, and connectionist techniques are instances of this general strategy. These methods are inductive in the sense that they move beyond their input and generalize it to create knowledge that was not previously in memory. In contrast, analytic learning methods formulate knowledge based mainly on other knowledge already in memory. Explanation-based learning is a clear example of this type of approach; it uses prior knowledge to explain new experiences, then simplifies the explanation and stores it in memory.

Inductive methods have also been called *knowledge-level* learners, as their acquired knowledge structures change the deductive closure of the system (Dietterich, 1987) (see KNOWLEDGE LEVEL). In contrast, analytic methods are *symbol-level* learners; their compiled knowledge may increase efficiency but leave the deductive closure unchanged. However, later some applications in which inductive techniques produce symbol-level learning and some conditions under which analytic methods can produce knowledge-level shifts will be considered.

Of course, more unified approaches are possible. A learner could use prior knowledge from the domain to constrain induction, producing partial explanations that still require inductive leaps to form general structures. Alternatively, instead of using deductive rules to construct explanations, a learner could use plausible inference rules; the resulting explanations may be plausible, but they are not guaranteed to hold, and thus involve a sophisticated form of induction. Some researchers have argued that the unification of inductive and analytic approaches should be given high priority (Langley, 1989).

FIVE PARADIGMS

This section will examine some recent research on machine learning. The field has been divided into five paradigms based on the basic representations and learning methods employed: connectionist approaches, genetic algorithms, empirical rule learning, analytic learning, and case-based methods. In each case, the basic approach is reviewed, recent advances are examined, and some open research issues are considered.

Each paradigm is loosely defined not only by a shared set of assumptions, concerns, and methods, but also by the amount of interaction among researchers. Generally, people within a given paradigm read and reference each other's papers, attend the same meetings and presentations, and all too often ignore work in other paradigms. Some important relations among these paradigms that reveal more similarities than appear at first glance are identified after discussing the five paradigms.

Connectionist (Neural Network) Learning

Some of the earliest research on machine learning focused on connectionist methods (Nilsson, 1965). Recently there

has been a resurgence of interest in this approach. The name derives from the basic representation for learned knowledge: a network composed of nodes connected by directed, weighted links (sometimes called neural networks (qv) because of the suggestive similarity between the computational style of network nodes and neurons). Such systems typically assume that inputs are represented as a set of binary features, with each feature being present or absent. In operation, features that are present activate the network's initial nodes. Then, the weights on links from these nodes to others determine whether subsequent nodes will be activated. The process iterates until activation has a chance to reach the network's final nodes, with the output of the network being the activation of the final nodes.

Let us examine the connectionist framework's position on the dimensions discussed in the previous section. Learning consists of modifying link weights to better mimic the desired relations among the inputs and outputs. Thus connectionist approaches lend themselves naturally to performance tasks that involve classification, but they can be adapted to more complex domains. Although not logical in form, the connections encode the acquired knowledge and summarize the data encountered. Researchers have focused predominantly on methods for supervised learning situations, but there are some exceptions. Connectionist researchers have also examined many alternative strategies for adjusting link weights; many are incremental, and all are inductive in nature.

Perceptrons and Linear Threshold Units. The simplest form of connectionist network is the perceptron or linear threshold unit (Rosenblatt, 1962). In this framework, there is a single output node to which each input node is connected by a single weighted link. The output node also has an associated threshold. Given a datum, the output node sums the weights of the links from active input nodes (those whose features are present). If this sum exceeds the threshold, the output node is activated; otherwise it remains inactive. Figure 1 presents an example of a simple perceptron.

Despite their simplicity, perceptrons can represent a variety of functions. For example, consider a network in which N links have a weight of one and all others are zero. If the threshold is set to N , the network encodes a rule that matches only when the conjunction of the nonzero features are present. Similarly, if the threshold is set to one, the same network encodes a disjunctive rule. Using a threshold of K (where $N > K > 1$) lets the network concisely express the K of N function. Unlike the former two, this latter function is difficult to represent using logical notation. Moreover, allowing weights other than one supports an even broader class of functions. In fact, given the appropriate weights, a perceptron can represent any linearly separable class. In other words, if we view the F features as defining an F -dimensional space, the network can describe any class that involves placing a single hyperplane between the instances of two classes.

There are a number of straightforward methods for learning appropriate link weights given example input-output pairs. One of these, the perceptron learning rule,

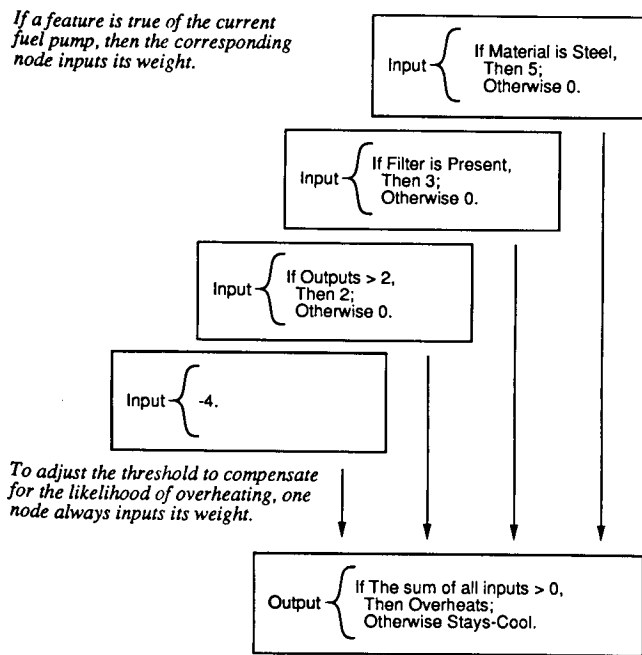


Figure 1. Schematic of the organization of a linear threshold unit. Inputs correspond to instance features (shown at top). If these features are present, their weights are summed by the output unit (at bottom). Learning modifies the weights associated with inputs and thereby alters the summed output. An additional constant input (lowest and leftmost input) allows the linear threshold unit to adjust its threshold.

learns only when it makes a prediction error. If the output unit is not active when it should be, then its incoming weights are too low; accordingly, they are incremented by a small constant. By the same reasoning, if the output unit is active when it should not be, the weights are too high and are decremented by the same small constant. Part of the appeal of this incremental method is that it is guaranteed to converge on any linearly separable class given a finite number of instances (Minsky and Papert, 1969).

Other learning methods, such as the LMS procedure (Widrow and Hoff, 1960), modify each weight differentially in an attempt to reduce the mean-squared error between the desired and generated output. This approach can be run incrementally or nonincrementally (using all available instances), and it can be generalized to handle networks with continuous rather than binary inputs. Linear threshold units have been used to learn regular and irregular endings for the past tense of English verbs (Rumelhart and McClelland, 1986), and a perceptron-like learning method has demonstrated its ability to acquire expert-level performance in the game of checkers (S Samuel, 1959).

Backpropagation in Multilayer Networks. Despite their attractiveness, single linear threshold units cannot represent or learn class descriptions that are not linearly separable. Although higher-order components such as quadratic threshold nodes might be feasible, a more common approach uses a network with multiple layers. These

structures include intermediate, or hidden, nodes that are indirectly connected to the inputs and outputs of the network. Given sufficient breadth in this structure, the network can express an arbitrary function of the input features. Moreover, if the structure is relatively shallow (ie, only short paths from inputs to outputs), large networks of this sort are computationally efficient to use. However, learning the weights for such networks is another matter entirely, and much of the recent research on connectionist learning has focused on this issue.

One of the most intuitive approaches to learning appropriate link weights in a multiple layer network applies the LMS procedure recursively. Known as *backpropagation*, this procedure first propagates activation through the network in the normal, forward direction. Based on the differences between observed and desired outputs, backpropagation uses LMS to compute the desired activation levels on the hidden nodes one level back. Not only does this indicate the appropriate weight change in the final links, it also allows backpropagation to treat the hidden nodes as if they were output nodes. Using the difference between the observed and desired activation for hidden nodes, backpropagation applies LMS recursively until it reaches the input nodes.

Like the perceptron learning rule, backpropagation is conducting a search in the space of link weights. Effectively, both carry out a hill-climbing search in which the gradient is defined by error reduction. Unlike the convergence result for the simpler learning rule, backpropagation may become stuck in local optima. This has not emerged as a significant problem in studies to date, but it remains an open issue. More pressing problems include a very slow rate of learning and some dependence on the number of hidden nodes. One noteworthy application demonstrates that backpropagation can acquire pronunciation knowledge that accurately predicts phenomes from English text (Sejnowski and Rosenberg, 1987).

Alternative Approaches to Connectionist Learning. Researchers have explored a number of other approaches to learning in multilayer networks. For example, Boltzmann machines (qv) are a probabilistic technique based on an analogy with thermodynamics in which nodes settle into stable configurations as the "temperature" of the system decreases (Ackley and co-workers, 1987). Such methods must be run many times, in order to reach equilibrium and to collect statistics about the probability of connected nodes being active simultaneously. As a consequence, they are typically even slower than backpropagation. Nevertheless, Boltzmann machines also have advantages, and active research continues in the area, including applications to speech recognition (Prager and co-workers, 1986).

Other researchers have taken a reinforcement-learning approach. In this framework, instead of fine-grained feedback about each of the network's outputs, the only information available is a single evaluation score for the network's overall behavior on each instance. For example, the AR-P algorithm rewards or penalizes each weight in the network equally as a function of the reinforcement evoked by the network's overall behavior (Barto and co-

workers, 1983). This learning scheme has been successfully applied to a number of domains, including a pole-balancing task involving the dynamic control of forces over time. Like the other methods, this approach extends to domains in which the inputs are real-valued rather than binary.

In other studies, researchers have explored the behavior and capabilities of cyclic and deeply nested network structures. For example, a cyclic network has its input nodes connected to its output nodes and vice-versa. This design allows networks to memorize patterns, and given a partial or noisy pattern to recall, the net can reconstruct the complete, noise-free original (Kohonen and co-workers, 1981). In a deeply nested network, hidden nodes are ordered, and each successive hidden node receives input from all prior hidden nodes. Nesting the hidden nodes capitalizes on the representations they have learned, and the resulting structure tends to generalize better (Fahlman and Lebiere, 1990).

Open Issues in Connectionist Learning. Research on learning in the connectionist framework has made significant strides since the early results with perceptrons, and initial applications have started to emerge. Still, a number of serious issues remain to be addressed.

- *Increasing the Rate of Learning.* Existing methods learn very slowly, often requiring many iterations through the training instances. Future research should examine the factors that affect the learning rate and develop connectionist methods that learn more rapidly (Fahlman, 1988; Hampson and Volper, 1987).
- *Generalization.* Current methods converge on weight settings that summarize training instances, but sometimes this is at the expense of accurate generalization over unseen instances. This issue becomes especially important when the training data are noisy (Fisher and McKusick, 1989; Knight, 1989; Weiss and Kapouleas, 1989).
- *Structural Knowledge.* Some domains seem inherently relational, but connectionist methods rely on feature-based representations. The framework must be adapted to represent and learn from relational and structural input (Hinton, 1986).
- *Sequential Behavior.* Connectionist techniques lend themselves to parallel implementations, but they have difficulty carrying out ordered actions like those required for problem solving. Extended architectures are needed that can handle sequential behavior, along with learning methods that can support them (Elman, 1990; Mozer and Bachrach, in press).
- *Incorporating Domain-Specific Bias.* Perceptrons have a strong bias toward learning linearly separable classes, but multilayer networks search a much larger space. Future research should examine methods for incorporating biases into multilayer networks that constrain the learning methods' search and improve their learning rates (Towell and co-workers, 1990).

Growing numbers of researchers are examining the connectionist paradigm seriously and many are concerned with issues of learning. Theoretical analyses and experimental studies have begun to reveal a deeper understanding of these methods' advantages and drawbacks, pointing the way to the extensions and improvements outlined above.

Genetic Algorithms and Classifier Systems

Genetic algorithms are a family of adaptive search methods that derive their name from a loose analogy with genetic change in a population of individuals. Like connectionist methods, most genetic algorithms assume a feature-based representation of instances and events. However, rather than using a weighted network to represent acquired knowledge, they employ a disjunctive set of knowledge structures or patterns. Each pattern is conjunctive and specifies the presence (or absence) of some features. Patterns also have an associated weight, sometimes called the pattern's fitness, that summarizes its performance on past experiences. Given a new instance, a stochastic scheme follows the recommendations of strong matching patterns to reach a decision (Wilson, 1987).

In terms of the dimensions of the preceding section, this paradigm lends itself naturally to performance tasks that involve both classification and problem solving, as is discussed below. Research on genetic algorithms has focused on supervised learning situations, although often the feedback is some overall score, as in the reinforcement learning framework discussed above. Researchers have explored a number of inductive approaches to incrementally modifying the content of individual patterns and their weights, and they have successfully applied genetic algorithms and classifier systems to a variety of tasks (Goldberg, 1990).

Learning with Genetic Algorithms. In addition to using a different learned representation, genetic algorithms and connectionist techniques also carry out a different type of search during learning. Whereas connectionist methods apply a hill-climbing operator to a single state (one set of weights), genetic algorithms apply several heuristic operators to a set of states (patterns). As shown in Figure 2, genetic algorithms follow three steps in processing new experiences: updating pattern strengths, applying search operators, and pruning ineffective patterns. Methods for updating weights vary widely among specific genetic algorithms; several are discussed below. For the second step, genetic algorithms typically use two operators, crossover and mutation, that apply to strongly weighted patterns to produce syntactically similar new patterns; crossover is analogous to gene splicing, whereas mutation introduces random, minor variations. To maintain the size of the pattern set, the third step replaces prior, weakly weighted patterns with new patterns. In a sense, patterns compete with each other to produce offspring in the next cycle or generation. There is a large body of both theoretical and empirical evidence showing that, even for very large and complex search spaces, genetic algorithms can rapidly lo-

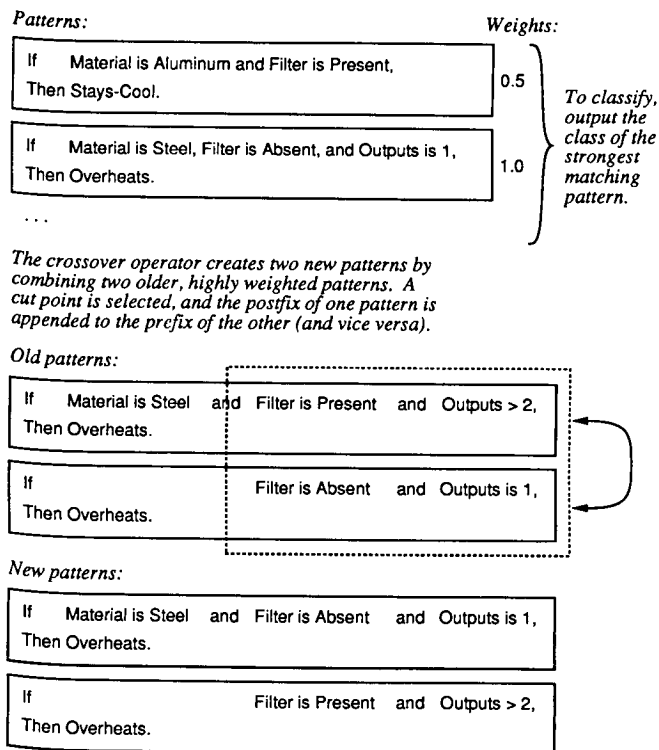


Figure 2. Schematic of the representations and operators used in the genetic algorithm. The weights of conjunctive patterns (shown at top) are used to classify instances. Learning modifies both the weights of patterns and their conditions. The simplest operator generates new patterns that are minor random variations of old ones (not shown). The crossover operator (shown at bottom) splices together an arbitrary prefix of one pattern with an arbitrary suffix of another (and vice versa).

cate effective knowledge structures using about 50 to 100 patterns.

Applying genetic algorithms to classification tasks is straightforward when instances can be represented as features. The set of knowledge structures is initialized to N random patterns, some of which will be very specific and others of which will be quite general. Each pattern is assigned a class. To adjust the weights, each time a pattern is successful (ie, it matches an instance from its own class or it fails to match an instance from another class), its weight is incremented; each time a pattern is unsuccessful (ie, it fails to match its class or matches another class), its weight is decremented. Because pattern weights are decremented when they fail to match instances in their class, if the concept to be learned is disjunctive, the weight of every pattern will be decremented at one time or another. Nevertheless, pattern strength is still a useful heuristic for directing search through the space of descriptions, and genetic algorithms can learn complex disjunctive concepts, even in noisy domains (Wilson, 1987). This ability results in part from the inclusion of many patterns in each generation, some of which compete and others of which complement each other, in that they come to occupy different niches.

Classifier Systems and Problem Solving. Classifier systems are an architecture for problem solving that incorporates a genetic algorithm as a component. Most work along these lines assumes a simple, forward-chaining production system, consisting of a set of condition-action rules and a dynamic working memory. Each rule is weighted and has one or more feature patterns as its condition and a single pattern as its action. Memory contains a set of fully specified patterns, or messages, that have a value for all features. Through encoding conventions, some messages originate as inputs to the system and others trigger outputs.

The classifier system framework extends the standard forward-chaining recognize-act cycle. On every cycle, each rule whose conditions match messages in memory makes a bid proportional to its weight. One or more rules are then selected for application, with probabilities proportional to their bids. The selected rules are applied by adding the patterns in their actions to working memory. These new messages may allow other rules to match, and the cycle continues.

There are two main aspects to learning in classifier systems. First, a genetic algorithm is used to generate new candidate rules from existing, strongly weighted rules (similar to the process described above). Second, the classifier system adjusts the weights of rules based on their contribution to desirable behavior. This involves assigning credit to useful rules and blame to faulty ones. An effective approach to this latter learning problem is called the *bucket brigade* (Holland, 1985). As rules apply, they pass along a portion of their weight to the rules that applied in the cycle before them. Some rule ultimately is rewarded directly by the environment, and this reward is iteratively passed back through the rules in the application chain, increasing their weights. In the simplest case, the weight of any rule that participates in the chain eventually converges on the amount received by the last rule in the chain. As a consequence, rules that apply but do not lead to external reward are consistently paying out their weight without receiving any reinforcement, and their weight diminishes. Classifier systems have been successfully applied to a variety of domains, including regulation of gas flow through pipelines (Goldberg, 1985) and survival in a resource-scarce environment (Booker, 1988).

Alternative Uses of Genetic Algorithms. Unlike classifier systems, which apply operators to propose individual new rules, another approach to using genetic algorithms for problem solving applies genetic operators to entire rule sets. Rather than exploring variations of rule conditions, the operators primarily explore variations on rule combinations. Because of this, the knowledge structure is composed of multiple rule sets, each of which constitutes a forward-chaining production system that has an associated weight.

In this framework, the weight of each rule set is evaluated by running the rules on a set of training problems. Because rule sets with strong weights tend to be selected by the search operators, useful combinations of rules are propagated through the knowledge structure, and less useful rule combinations are gradually eliminated. The

operators occasionally introduce new rules, but these are always evaluated in the context of their rule set. The power of these ideas has been demonstrated by a state-of-the-art poker-playing system (Smith, 1983) and by an effective system for multiple class discrimination in the domain of human gait analysis (Schaffer and Grefenstette, 1985). Recent research also shows that a combination of this approach and classifier systems performs better in some domains than either in isolation (Grefenstette and co-workers, 1990).

Open Issues in Genetic Algorithms. Although considerable progress has occurred in the understanding of genetic algorithms since their inception, many open research issues remain. These include:

- *Alternative Representations.* Genetic algorithms typically assume a feature-based representation of knowledge. Future work should explore the application of these methods to more sophisticated representations (Gordon and Grefenstette, 1990; Koza, 1989), adding new genetic operators if necessary.
- *Acquired Representations.* Classifier systems can represent complex class descriptions by the organization, variability, and distribution of weights in clusters of rules. However, we need to better understand the underlying nature of such learned representations and how are they acquired (Belew and Forrest, 1988).
- *Emergence of Useful Symbols.* Classifier systems can use tags to build associations between rules, producing behavioral sequences. Recent work has explored the development of such internal symbols and the conditions under which they emerge (Shaefer, 1987).
- *Credit Assignment.* The issue of credit assignment is central to applying genetic algorithms to problem solving, and a variety of methods have been proposed, including weight update, conflict resolution, and the use of "taxes." Research is needed to determine the conditions under which each approach behaves well and to explore hybrids that might do better than any method in isolation (Grefenstette and co-workers, 1990).
- *Incorporating Domain Knowledge.* Genetic algorithms seem especially well suited for knowledgeable domains in which extensive search is necessary, but they may also be able to use and refine existing domain knowledge (qv).
- *Population Size.* In some domains, genetic algorithms should behave significantly better than hill-climbing techniques (eg, connectionist methods). Some research has studied the effect of the number of patterns in a knowledge structure (Robertson, 1988), but future work should identify the broader conditions under which maintaining a redundant knowledge structure is worth the cost.

Researchers in the genetic algorithm community are already attacking these problems, but more work remains before this promising approach achieves its full potential.

Empirical Learning Methods

Another community of machine-learning researchers have studied empirical methods for the acquisition of more explicit knowledge structures. Let us consider the dimensions of machine-learning research for this approach. Like connectionist and genetic methods, these techniques are inductive, in that they move beyond training instances to make predictions about novel cases. Unlike them, empirical learning methods have employed relational and structural representations for both experiences and acquired knowledge, though the majority of research has thus far focused on propositional representations. Much of the early work on empirical learning dealt with classification domains, but there has also been progress on problem solving and natural language acquisition.

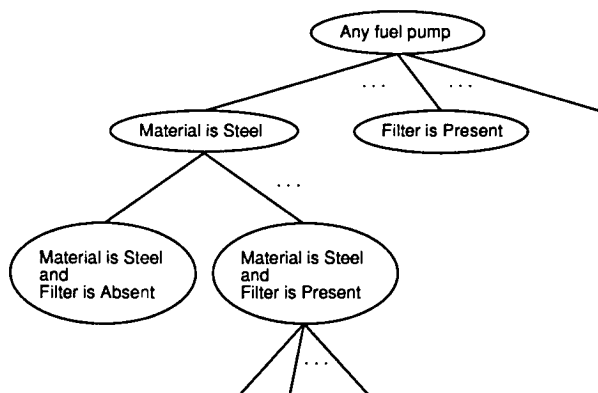
Currently there are a variety of well-understood methods that learn production rules, decision trees, and concept hierarchies (to name a few of the many learned representations). Some methods require the close supervision of a tutor, as described in the above section on Supervised and Unsupervised Learning, whereas others learn in an unsupervised fashion. Some techniques require all instances at the outset, whereas others learn incrementally and can process new instances with little additional effort. This section reviews three main approaches to empirical learning, along with some open research problems.

Empirically Learning Production Rules. One common scheme for representing domain expertise uses production rules whose conditions test properties of experiences and whose actions specify classifications. Researchers have explored a variety of empirical methods for learning such rules from a set of preclassified training instances, and most approaches rely on the fact that the space of rule conditions are partially ordered according to generality. Thus, one can start with the most specific possible description, using a generalization operator to remove or relax conditions. Alternatively, one can start with the most general possible description, using a specialization operator to add or constrain conditions. The candidate-elimination algorithm (Mitchell, 1977) employs both these ideas to carry out a bidirectional exhaustive search to identify conditions for classification rules. For each possible class, the algorithm maintains a version space that summarizes the space of hypothesized conditions in terms of a most-specific boundary set and a most-general boundary set. New positive instance may indicate the need for more general descriptions, forcing revision of the specific boundary, whereas negative instances may suggest more specific descriptions, leading to revision of the general boundary. This continues until the algorithm converges on a single conjunctive description in both sets, or until one of the boundary sets becomes empty, indicating an inconsistency.

The candidate-elimination algorithm assumes that a single, conjunctive rule can describe each class, and that training instances are free of noise. However, disjuncts and noise are common in applied settings. Another appealing family of learning methods relaxes these assumptions and uses heuristic search to limit computational ex-

pense. These methods employ beam search (see SEARCH, BEAM) or related methods to find individual rules that discriminate between positive and negative instances of a class. Search may occur from general to specific rules or in the opposite direction. During each search, candidate rule conditions are minimally specialized (or generalized) in all possible ways, each specialization (or generalization) is heuristically evaluated for predictive accuracy on the training instances, and the best are further modified. Search terminates when none of the new specializations (or generalizations) are statistically better predictors than their predecessors, at which point the best candidate is used to construct a classification rule. To handle disjunctive domains, some methods then remove all positive instances from the training set that are covered by this rule and repeat the search process over the remaining instances, continuing until all positive instances are covered by some rule (Michalski, 1983; Clark and Niblett, 1989). Figure 3 depicts this approach graphically.

Methods of this sort have been successfully applied to moderately realistic tasks. For example, in the domain of lymphography, some rule-learning systems (Michalski, 1987; Clark and Niblett, 1989) have equaled the classification accuracy of human experts (82% correct). Another rule-learning system (Schlimmer, 1987) has demon-



Search moves from general rule conditions to specific ones until the matching pumps are not statistically distinguishable.

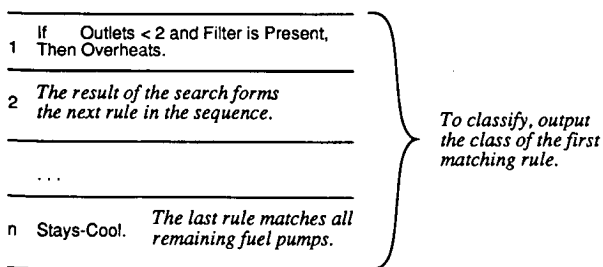


Figure 3. Schematic of the representation and search used in CN2, a recent rule learning method. Learned rules (shown at bottom) are ordered, with the last rule identifying the class of any remaining instances. Search maintains a limited size boundary as it looks for rule conditions that are both predictive and reliable. Initially, it looks for a consistent description of a single class given all the instances, and subsequently it searches for a consistent description of those examples not covered by any rules found thus far. Search terminates when all instances are covered by some rule.

strated similar results for recognizing poisonous mushrooms (95%) and predicting the political party of Congressmen based on their voting records (90%).

In spirit, methods for learning production rules are similar to connectionist and genetic methods. However, there are some significant differences. As we have seen, most of these methods take advantage of the fact that the space of class descriptions can be partially ordered by generality. Thus, most rule-learning schemes search using specialization operators, generalization operators, or both. Second, the search in most rule-learning methods is biased toward finding simple descriptions. If this is appropriate in a given domain, then these methods tend to learn much more quickly than connectionist or genetic methods. Finally, some rule-learning methods represent experience and learned knowledge using relational and structural representations, giving them more expressive power than other approaches.

Although methods for empirical rule learning were originally designed with classification domains in mind, they can also be applied to problem-solving tasks. For example, given a set of legal operators for carrying out state-space search, the same methods can acquire the heuristic conditions under which each operator should be applied. However, this approach requires first identifying appropriate and inappropriate applications of each operator, and this is equivalent to assigning credit and blame to steps along a search path. One response to this latter issue involves waiting until a complete solution has been found. Then, steps along the solution are labeled as appropriate operator applications, whereas all steps leading off the solution are labeled as inappropriate (Langley, 1985; Mitchell and co-workers, 1983). Another response to the credit/blame assignment issue involves interacting directly with a domain expert who provides immediate feedback about the desirability of each action. In either case, one can then apply empirical learning methods to the appropriate and inappropriate applications (as positive and negative instances of the class, respectively), producing heuristic rules as output.

Constructing Decision Trees. Other work on empirical learning takes quite a different approach to the supervised learning task. This framework assumes the same input as systems that learn production rules (ie, a set of instances assigned to classes), but the learned knowledge is represented as a decision tree (Brieman and co-workers, 1984; Quinlan, 1983). Each nonterminal node of this tree specifies some attribute to test, each branch specifies an alternative value, and each terminal node specifies a class. To classify a new instance, a decision tree iteratively tests nonterminal node attributes of that instance and follows matching branches until it reaches a terminal leaf that classifies the instance.

The most common decision-tree learning method uses a divide-and-conquer algorithm, selecting domain attributes to partition the instances and recursively building subdecision trees to describe partitions. An evaluation function selects the most discriminating attribute for each nonterminal node's test. Instances are partitioned based on their value for the test attribute, and subtrees are con-

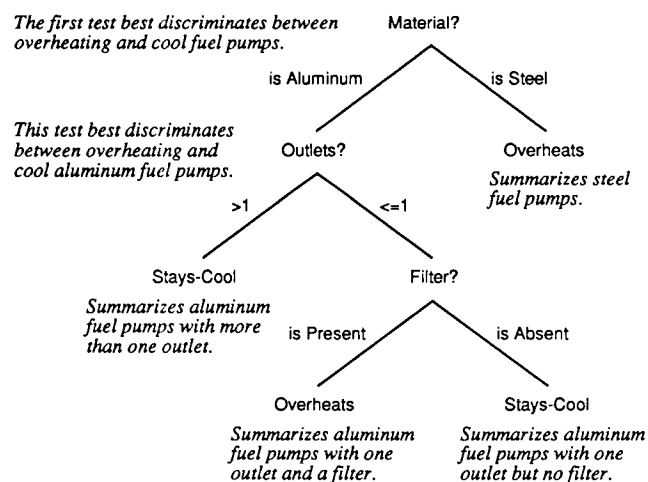


Figure 4. Schematic of the representation and search used in decision tree learning. To classify an instance, the test at the root of the tree (shown at top) is applied first, and based on the outcome, classification proceeds with the appropriate subtree. The search begins by determining the single most discriminating test. Instances are then partitioned by their outcome for this test. The standard method terminates search when all instances in a partition are of the same class (or there are no more tests), but some variants use statistical tests to halt earlier as an effort to avoid overfitting.

structured to describe each partition. As the process iterates, subtrees are complete when all instances in their partition have the same class (or when there are no more attributes to test). This process can be viewed as a greedy, general-to-specific search through the space of decision trees. Figure 4 shows an example of this approach.

Decision-tree methods have been applied to a variety of classification tasks, representing a mix of synthetic and natural domains. For example, induced decision trees can successfully recognize lost chess endgames (100% for losses in three-ply) (Quinlan, 1983), and they can accurately classify thyroid diseases (99%) (Quinlan and co-workers, 1986). The latter application resulted in a decision tree that outperformed a handcrafted expert system that took years to construct.

Methods for inducing decision trees constitute some of the most widely studied algorithms within the machine learning community. Extensions include techniques for pruning trees in response to noisy training data (Brieman and co-workers, 1984; Kononenko and co-workers, 1984; Quinlan, 1986b), methods for detecting useful thresholds on numeric attributes (Quinlan, 1986a), and algorithms for incrementally revising a tree in response to new data (Schlimmer and Fisher, 1986; Utgoff, 1989). Decision-tree methods have received significant attention within applied AI circles, and some industrial groups have used them to construct expert-level diagnostic systems automatically.

Forming Concept Hierarchies. In unsupervised concept learning tasks, no expert is available to classify instances for the learner. Instead, the learner is given a set of unlabeled instances and is asked to form "useful" concept de-

scriptions. A common approach to this problem, called *conceptual clustering* (see CLUSTERING) (Michalski and Stepp, 1983), involves (1) determining how to cluster the instances and (2) building descriptions for those clusters. Typically, conceptual clustering methods form a hierarchy or taxonomy of concepts. Although they are superficially similar to decision trees, each node in a concept hierarchy has an associated concept description that is used during classification. Also, although both paradigms typically involve a top-down search from simple to more complex representations, conceptual clustering methods use different search operators and evaluation functions than those used for building decision trees. For example, some techniques measure the simplicity of the potential clusters' descriptions; others aim to maximize predictive accuracy over all attributes, since no class information is available.

Conceptual clustering is an active area of research within the empirical learning paradigm. Recent work has focused on using goals to direct the clustering process (Stepp and Michalski, 1986), on devising incremental methods for hierarchy formation (Fisher, 1987; Lebowitz, 1987), and on the role of probabilistic descriptions in clustering (Cheeseman and co-workers, 1988; Fisher, 1987). These methods hold considerable promise because they address the issue of organizing concepts into integrated memory structures.

Methods for conceptual clustering appear to be appropriate for classification domains where no expert is readily available. For example, musicologists viewed the taxonomy formed by one conceptual clustering algorithm as a significant scientific contribution (Michalski and Stepp, 1983). Similarly, a learned set of classes of infrared stellar bodies were judged significant by astronomers (Cheeseman and co-workers, 1988). Finally, conceptual clustering methods have led to improved performance on prediction tasks for other domains, including soybean diagnosis and congressional voting records (Fisher, 1987).

Open Issues in Empirical Learning. Although our understanding of empirical learning methods has reached the stage where initial applications are feasible, basic research is still in progress. Active areas of research include:

- *Incremental Learning.* Many existing empirical methods are nonincremental: They must reprocess many instances to incorporate new information. Incremental methods can be more efficient, but most are not yet as robust as their nonincremental counterparts (Fisher, 1989; Utgoff, 1989).
- *Search-Limited Methods.* Many existing empirical techniques carry out a significant search through their space of descriptions. Search-limited methods, such as hill climbing and greedy algorithms, are much more efficient, but they should be modified to increase their chances of finding satisfactory solutions (Quinlan, 1986a; Langley and co-workers, 1987).
- *Incorporating Domain Knowledge.* Most empirical learning systems use domain knowledge in minimal

ways. These methods should be extended to use available knowledge to constrain search and produce clearer knowledge structures (Drastal and co-workers, 1989; Elio and Watanabe, in press; Hirsh, 1989).

- *Representation Change.* Most learning systems are unable to extend their initial representation language. Empirical methods for defining new terms hold great promise, but they must constrain their search for such terms and generate effective candidates (Muggleton and Buntine, 1988; Pagallo, 1989; Matheus and Rendell, 1989).
- *Noise and Concept Drift.* Some methods are robust with respect to noise and changing environments (Schlimmer and Granger, 1986), but researchers should explore the general principles underlying these issues and identify the class of techniques that can handle them.
- *Uncertainty and Probability.* Many expert systems for classification employ probability or other techniques for handling uncertainty, but few learning methods incorporate these ideas. Researchers need to extend existing representations and methods so as to handle uncertainty (Fisher, 1987; Geiger and co-workers, 1990).

Research on empirical techniques is leading to continual progress on the foregoing problems, but much remains to be done in this promising area of machine learning.

Analytic Learning

In contrast to the inductive learning methods discussed so far, another major paradigm in machine learning focuses on analytic learning. This approach emphasizes the transformation of existing domain knowledge into a more useful form, using data only to guide the application of deductive processes to this knowledge. These methods are sometimes called *symbol-level learners* (Dietterich, 1987) because the learned knowledge increases efficiency. [However, in some applications, inductive methods can also improve efficiency (Langley, 1985; Mitchell and co-workers, 1983).] If the performance system must operate under limited resources, such as time or memory, then such learning can indirectly lead to improvements in accuracy as well.

In terms of the dimensions of machine-learning research, learning methods in this paradigm typically encode experiences, domain knowledge, and learned knowledge with relational representations. Because efficiency is the most obvious benefit of analytic learning, these methods have typically been applied to problem-solving performance tasks. They have also been studied in both supervised and unsupervised learning situations; in addition to constructing new operators, these methods have also been applied to the task of improving operator selection. As many have observed, methods of this type could be applied to existing knowledge without the aid of data to guide learning. This nonincremental approach has been predominantly shunned in favor of more computationally efficient, incremental approaches, though there are recent exceptions (Etzioni, 1990).

Explanation-based learning is one common approach that can be viewed as compiling knowledge into an efficient form rather than creating or extending knowledge. This class of learning methods can simplify a problem solver's reasoning process by composing rules into useful combinations. Explanation-based learning can also acquire control knowledge that limits alternatives, thus reducing the amount of search problem solving incurs. There currently exists a number of well-specified algorithms for explanation-based learning, and recently there have been some successful applications to significant problem-solving tasks. This section reviews the basic approach, its adaptation to problem solving, and open research questions.

Compiling Explanations into Rules. Much of AI research on reasoning and theorem proving takes a problem-reduction approach. This framework assumes that domain knowledge is specified as a set of inference rules or goal decompositions. For example, to get from Washington to New York, an agent can drive to National airport, fly to LaGuardia, and take a taxi into the city. Given a top-level goal, an agent can use AND-OR search to find some set of primitive actions, states, or beliefs that will achieve that goal. Some programming languages, such as PROLOG, support this form of reasoning directly.

The result of a problem-reduction search is a proof tree or explanation for how to achieve the initial goal. Methods for explanation-based learning use this information during the learning process to create summarizations of search that can simplify future search for similar goals (DeJong and Mooney, 1986; Mitchell and co-workers, 1986). More precisely, given a problem and its solution, explanation-based learning uses a proof tree for the solution of a problem to: (1) focus attention on relevant problem features, and (2) summarize the problem-solution pair as a general rule. The resulting rule states the conditions under which the proof will hold, and, in the future, similar problems can be solved in fewer search steps.

Figure 5 summarizes the basic ideas that underlie explanation-based learning. This approach proves an important demonstration of the use of knowledge in learning, and the basic method can be applied to any domain in which knowledge can be stated as monotonic production rules (ie, rules that only add knowledge). This includes problem-solving domains for which useful goal decompositions are already known, as well as many reasoning tasks and design problems.

Learning Macro-Operators. Not all research on problem solving views this process in terms of AND-OR search. In contrast, some work focuses on state-space search, in which one applies a sequence of operators to problem states in order to achieve some desired state or goal. For example, in the blocks-world domain, states and goals involve specific configurations of objects, whereas operators specify the preconditions and results of actions that manipulate objects. Explanation-based methods can use knowledge of such legal operators to construct general rules that increase problem-solving efficiency (Fikes and co-workers, 1972).

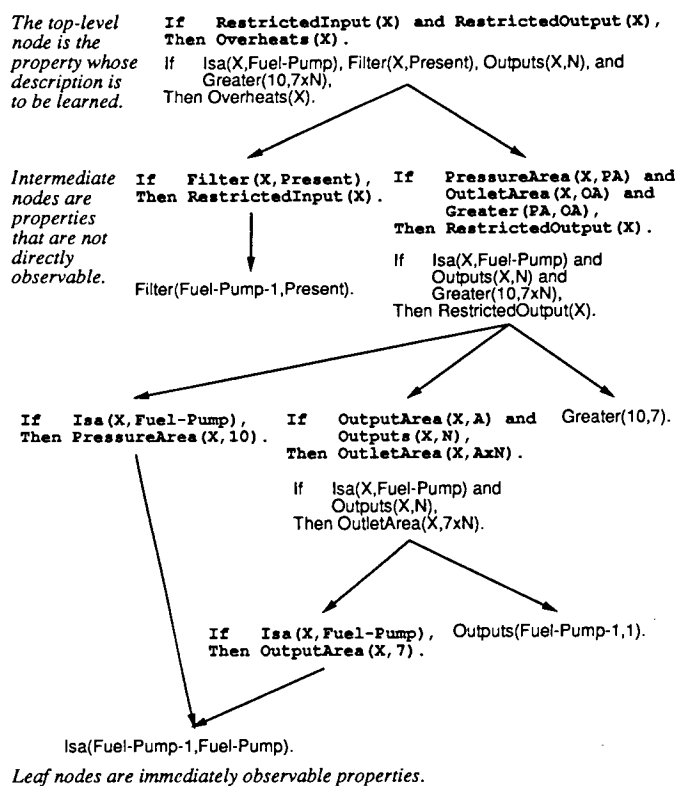


Figure 5. Schematic of the representation used by explanation-based generalization. Properties of an instance (depicted as leaves) and domain knowledge encoded as rules (depicted in teletype font) justify or *explain* other properties of the instance that may not be immediately observable. By extracting the weakest preconditions under which this justification holds, one can construct rules that are more general than the specific instance but more special than the general domain knowledge rules (depicted below each domain knowledge rule in sans-serif font). These rules allow inferring a non-observable property by testing only immediately observable properties of an instance.

The most basic application of explanation-based learning to state-space search is straightforward. Once a problem-solving system has found a sequence of operators that transform an initial state to a goal, one composes this solution path into a single rule or macro-operator (Fikes and co-workers, 1972; Iba, 1989). The conditions of this rule include all aspects of the initial problem state that were required for the solution to hold, and the results include all those actions not undone by others along the way. This composition process is more complex than the one just described because operators can be nonmonotonic (ie, they add and delete facts). However, it is simpler in other ways because solution paths are sequential rather than tree structured.

The construction of macro-operators lets a problem solver take larger steps through a problem space and thus shorten the effective length of solution paths. In contrast to inductive learning methods, macro construction is a purely deductive process: the form of the new rule is completely determined by the primitive operators and their instantiations. Most work on discovering macro-operators has focused on abstract tasks such as the eight puzzle and

the blocks world, but it should apply to any task that can be cast in terms of state-space search. The same basic method also applies to planning approaches, in which sub-goals are created during the problem-solving process (Minton, 1985).

Learning Preference Rules for Problem Solving. Another approach to explanation-based learning generates heuristics for problem solving. Although early work on rule-based problem solving used domain-independent heuristics to select which state to expand and which operator to apply during search, some recent learning research has focused on the acquisition of domain-specific search control rules that conditionally prefer one state or operator to another. In the section on Empirically Learning Production Rules, it was shown that one could use empirical methods to learn such heuristics, but this task also fits the requirements of explanation-based methods. [In particular, preference rules are monotonic (ie, they only add facts), so they may be acquired using explanation-based methods.] More work on learning search control has occurred within this paradigm.

In one version of this approach, a mean-ends planning system uses control rules to select a state to expand, an operator to apply, and a binding for the variables in the operator (Minton and co-workers, 1989). When no control rules are available, the problem solver defaults to performing a depth-first search (see SEARCH, DEPTH-FIRST). When search leads to failure or success, the system tries to explain that failure or success using general knowledge of problem solving. In both cases, explanation-based learning compiles the AND-OR explanation into a control rule that the system stores away for future use. This approach has been tested on a number of planning domains, including moderately complex scheduling tasks.

In a related but somewhat different approach, new control rules are learned only when existing control rules result in an ambiguous decision or impasse (Laird and co-workers, 1986). In such cases, the problem-solving system searches to determine the correct answer and compiles the result into a new control rule. Instead of using explicit knowledge about the operators themselves, the conditions of the new control rule incorporate those facts that were used in determining the correct decision, while its actions contain the results of the search. This approach has been tested on a wide range of domains, including design, puzzle-solving, and a computer configuration task (Rosenbloom and co-workers, 1985).

Open Issues in Analytic Learning. Although there have been rapid advances in our understanding of analytic approaches to learning, there remain many significant research issues. The current open problems include:

- *Incorrect and Incomplete Knowledge.* Most existing analytic methods rely on a complete and correct knowledge, but this is seldom a realistic assumption. Researchers should develop methods that behave robustly when some knowledge is missing or faulty (Laird, 1988; Smith and co-workers, 1985).

- *Extending and Revising Knowledge.* Another response to incomplete and incorrect domain knowledge is to extend and revise the knowledge base (Carbonell and Gil, 1987; Muggleton and Buntine, 1988; Ourston and Mooney, 1990). We need methods to detect such problems and alter the knowledge base; note that this requires some form of inductive learning, though deduction may also play a role.
- *Intractable Knowledge.* For useful learning to occur in this framework, the system must already be able to perform search tractably, at least until sufficient knowledge is acquired to control this search. Knowledge bases for some domains, such as chess, are complete but intractable. What types of methods can handle approximate knowledge bases (Tadepalli, 1989), and how can this knowledge be used effectively? Some promising approaches include the use of abstraction (Ellman, 1988; Knoblock, 1990).
- *Nonlogical Knowledge.* Not all domain knowledge is logical in content; in some domains one must rely on heuristic rules or probabilistic relations. Researchers should extend explanation-based methods to use such knowledge to construct explanations and to generate useful rules.
- *Evaluating Alternative Explanations.* Given multiple proofs, a learning method must decide which to compile into the knowledge base for future use. Researchers need to identify the basis for judging an explanation's quality, and they need to devise efficient implementation techniques for this (Ng and Mooney, 1990).
- *The Utility Problem.* Having selected an explanation, a learning method must still determine whether the resulting compiled rule is worth retaining. Additional knowledge can increase retrieval costs and branching factors (qv) on future problems, and learning methods must decide if the cost outweighs the benefit. Recent advances have focused on simplifying compiled rules (Keller, 1987; Minton, 1990), limiting the expressiveness of acquired knowledge (Tambe and Rosenbloom, 1989), and collecting statistics on rule utility (Markovitch and Scott, 1989; Minton, 1990).

As with the other approaches to learning that have been examined, research on explanation-based methods continues to make steady progress, but much more remains to be done.

Case-Based Methods and Analogy

There is mounting evidence that human experts rely at least partly on memory for individual cases, particularly in domains such as law (qv), mathematics (see INDUCTION, MATHEMATICAL), design (qv), and planning (qv). Thus, it seems natural to exploit this idea in constructing AI systems, using memory of specific cases to classify new cases and to formulate plans. This is commonly called the case-based approach, and it constitutes a fifth major paradigm of machine-learning research (see REASONING, CASE-BASED).

Work on reasoning by analogy also falls within this general framework.

Unlike many other machine-learning methods, a general theme of case-based methods is that abstraction of prior experience primarily occurs in a lazy fashion. Rather than aggressively abstracting or compiling experience in anticipation of future use, case-based methods typically save the bulk of their processing until an actual use occurs for this experience. With this deferred processing scheme, three fundamental issues arise: (1) retrieving case(s) that may help with a new case, (2) matching and applying the retrieved case(s) to the new case, and (3) storing the outcome of the new case for future use. Rather than discuss each of these issues separately, this section reviews three subparadigms within the case-based paradigm and discusses the issues as they arise.

Consider the placement of case-based methods on the dimensions of machine-learning research. This framework represents individual experience either propositionally or with relational languages, although the latter are not as widespread because of the undesirable matching complexity that may result. Like empirical learning methods, case-based methods have been applied to both classification and problem-solving tasks; they have been applied in both supervised and unsupervised learning situations; and they are predominantly incremental. Current case-based methods always perform induction (usually at retrieval time), but many include an analytic learning component as well.

Nearest-Neighbor Techniques. In the simplest variant on the case-based framework, one simply stores past instances verbatim. When a new case is encountered, one finds the best match from among the stored instances, and then uses that case to supply the missing information directly. For example, in a medical diagnosis domain, each case would consist of a patient's symptoms along with his disease. Given a new patient's symptoms, this simple variant enumerates stored cases to find the "nearest neighbor" and uses its associated disease to predict the patient's malady. Figure 6 illustrates the basic technique. Nearest-neighbor algorithms are not limited to classification domains, as shown in recent work using case-based techniques for state-space problem solving (Bradtke and Lehnert, 1988).

Variants on this method are possible, such as expanding retrieval to making predictions based on a weighted average using the K nearest neighbors. Applying this technique to the task of mapping letters to phonemes achieves an 88% predictive accuracy on a test set of 1024 instances (Stanfill, 1987). The same research also demonstrates that this method degrades gracefully as one adds noise and as one decreases the size of the case base.

Uncontrolled growth of the case base is a natural concern in this paradigm. Thus, one may want to store cases selectively and delete others on occasion. In one instance of this approach, new cases are stored only when the existing knowledge base leads to a classification error (Aha and co-workers, in press). Although relatively unsophisticated, the predictive accuracy of this approach compares favorably to methods for inducing decision trees, and a




Figure 6. Schematic of the representation used in nearest neighbor learning methods. Instances are stored explicitly in a space of N dimensions, where each dimension describes one feature of the instance. The figure depicts the two-dimensional case. To classify an instance, one finds the nearest stored instance and predicts its class. It is possible to define a decision boundary representing equal distance from stored instances of different classes (shown as a curved line in the middle of the figure). Typically these decision boundaries are not stored explicitly.

similar approach has been successfully applied to the challenging domain of speech recognition (Bradshaw, 1987).

Analogical Matching. Given feature-based or attribute-based representations, the process of matching two cases is simple and inexpensive. However, domains like planning and design require structural or relational representations, and these introduce serious complexities into the match process. Because the new case and the stored case are unlikely to match exactly, one must perform some form of partial matching, and this problem is exponential (Watanabe and Rendell, 1990).

Researchers concerned with the process of analogy have devoted considerable attention to this issue, with most approaches involving some form of heuristic search through the space of partial matches. In this framework, the main issue becomes findings ways to constrain and direct the search for a useful match. For instance, one approach finds mappings that preserve higher-order relations between two cases in preference to ones that preserve simple features shared by the cases (Falkenhainer and co-workers, 1989). Other researchers (eg, Winston, 1984) have proposed different but related methods. In a complementary vein, knowledge about the domain can help evaluate potential matches and identify meaningful partial matches (Koton, 1988b).

Indexing and Memory Organization. Even with selective storage of cases, complex real-world domains might require thousands of instances, some having considerable structure that requires complex analogical matching. In such situations, one cannot afford to match against all cases stored in memory exhaustively, and the initial storage and subsequent retrieval of relevant cases become central issues. The natural response is to index cases by appropriate features, thus making the retrieval process more selective and reducing the effect of memory size. Early work on analogy focused on matching to the exclusion of indexing issues, but this has changed in recent years.

The first step in indexing cases involves selecting an appropriate set of indexes. The programmer can fix the indexes at the outset, but this produces an inflexible system that cannot adapt to new domains. Some researchers (eg, Lebowitz, 1987) have invoked inductive learning methods to identify predictive features, which are then used as indexes. Others have used explanation-based techniques to determine relevant features for each case and index on these instead. Problem-solving domains are especially well-suited to the latter approach, as the trace of problem-solving behavior (ie, goal trees) provides ready-made information for explaining success or failure (Carbonell, 1986; Hammond, 1986), which can then be used to index complex cases. The notion of derivational replay (Mostow, 1989) based on such problem-solving traces has received attention in many circles, including software engineering and automated VLSI design.

However, indexing by itself is not sufficient to allow efficient retrieval of relevant cases. For large knowledge bases, one must also organize memory into some manageable structure. Discrimination networks (Feigenbaum, 1963) are one approach to memory organization, but retrieval of a case often depends on a conjunction of features being present. This leads to fragility in domains where features can be missing, but the basic approach can be extended to support redundant indexing (Kolodner, 1983; Lebowitz, 1987; Fisher, 1987). In addition, one can store abstract summary descriptions at internal nodes in the network, giving generalization beyond individual cases. The construction of prototypes in this manner reveals an underlying similarity between case-based learning and conceptual clustering, which was discussed under Forming Concept Hierarchies.

Open Issues in Case-Based Learning. Research on case-based approaches has led to a number of promising methods, some of which have been tested on challenging domains. However, a number of open issues remain to be addressed:

- *Selecting Indexes.* A number of methods exist for selecting indexes. However, we need to identify common processes that underlie these approaches and the common information they exploit (Bareiss and co-workers, 1987; Hammond, 1986). We also need to better understand methods that generate indexes dynamically (Barletta and Mark, 1988; Kolodner, 1989; Owens, 1989).

- *Memory Organization.* Some initial work has addressed the organization of memory, but we need to identify the general properties that a memory should exhibit and to improve methods for dynamically reorganizing memory as new cases are encountered (Kolodner, 1983).
- *Matching Metrics.* Many existing techniques employ ad hoc schemes for matching against cases in memory. Researchers need to search for underlying principles involved in determining a good match and to develop methods for predicting good matches (Kolodner, 1989; Koton, 1988a; Salzberg, 1990).
- *Multiple Cases.* Analogy has predominantly focused on using sophisticated information from a single case, whereas K nearest-neighbor methods illustrate how to use simple information from multiple cases. Some early work has focused on combining highly relevant but contrasting cases (Ashley and Rissland, 1988), but we need to identify other types of information that multiple cases can provide (Aha and Kibler, 1990; Redmond, 1990).
- *Connections Among Cases.* In some domains, each case may have a complex, internal structure, effectively consisting of many component cases. Researchers should devise principled representational schemes that can capture connections between component cases and find methods for efficiently storing and retrieving such structures (Jones, 1989; Redmond, 1990; Sycara, 1988).
- *Forgetting.* Although most cases are useful, storage of all cases can lead to overfitting effects in noisy domains. We need techniques that can efficiently determine when to forget cases (Aha and co-workers, in press). An alternative is to develop methods that selectively acquire knowledge, thereby limiting the discovery of useless cases and unnecessary processing (Hunter, 1990).

Work on case-based reasoning has produced some promising techniques, but researchers need to explore the space of such methods more fully and to evaluate alternative approaches carefully in terms of their performance on real-world domains.

Relationships Among the Paradigms

Historically, machine-learning researchers have emphasized differences among the five paradigms we have discussed, rather than their similarities. This trend has been encouraged by differences in terminology, notation, test cases, and methods of evaluation. For instance, researchers studying connectionist techniques, genetic algorithms, and rule induction often run experiments with their inductive methods, but they typically use different data sets and measure different aspects of learning behavior. The same problem occurs between workers in the explanation-based and case-based paradigms. Our discussion so far has reflected this trend, focusing on the differences between the five learning frameworks.

However, understanding the similarities among these

paradigms is equally important to the science of machine learning. To this end, let us briefly consider some possible connections.

- *Symbolic and Subsymbolic Induction.* Many distinguish the subsymbolic approach of connectionist and genetic algorithms from the symbolic approach taken by empirical methods for rule induction and decision-tree construction. But despite differences in the representation of acquired knowledge, the spaces searched, and the learning operators employed, all three approaches are inductive in nature, and one can generally apply them to the same learning tasks. Recent comparative studies have clarified this fact (Mooney and co-workers, 1989; Dietterich and co-workers, 1990), producing comparable results for a variety of methods.
- *Induction and Explanation.* Researchers often make a dichotomy between inductive (eg, empirical) methods and analytic (eg, explanation-based) ones, characterizing the former as knowledge lean and the latter as knowledge intensive. Yet there is nothing mutually exclusive about these approaches, and hybrid methods should prove better than either in isolation. For instance, recent work on empirical methods has shown that domain knowledge and deduction can improve learning (Drastal and co-workers, 1989; Elio and Watanabe, in press). Similarly, one can use empirical methods to extend incomplete domain knowledge and to revise incorrect rules (Carbonell and Gil, 1987; Ourston and Mooney, 1990).
- *Explanations and Cases.* Explanation-based methods use domain knowledge to construct explanations and compile rules, but some case-based techniques rely just as heavily on domain expertise (Braverman and Wilensky, 1990; Redmond, 1989). Although one approach stores general rules and the other stores specific cases, the reasoning processes can be remarkably similar.
- *Cases and Abstractions.* Researchers often emphasize the distinction between storing specific cases and forming abstractions. Yet many case-based systems create abstractions as indexes for cases (Kolodner, 1983; Fisher, 1987), making them more accurately described as hybrids. Nor must a system that stores cases always use them during performance; in some situations, a hybrid system may prefer to use an abstraction. (Fisher, 1989). Moreover, work on conceptual clustering and on case-based learning shares a concern with the organization of memory, relying on similar structures and mechanisms.

In summary, there is considerable overlap between the paradigms in both their concerns and their approaches, although this is seldom apparent from research papers. Machine learning has just begun to converge on a set of standard terms and notations for describing systems and on a set of standard test beds and experimental methodologies for evaluating systems.

As researchers start to communicate across paradigm

boundaries, they can begin exploring the relationships more seriously. For instance, one can imagine a unified theory of induction that explains the behavior of decision-tree methods, genetic algorithms, and connectionist networks and that predicts the conditions under which each method would be most appropriate. One can also expect the development of hybrid algorithms that cut across paradigms to achieve better results than either in isolation. Some research along these lines has already begun, such as recent work on combining decision trees with perceptrons (Utgoff, 1988). We hope that some researchers will concentrate their efforts on such cross-paradigm research, as this may lead to new techniques that otherwise might never come to light. If such systems were successful, this would further strengthen the ties between areas, ultimately transforming machine learning into a unified field rather than many subdisciplines.

METHODOLOGICAL DEVELOPMENTS

Machine learning is a scientific discipline; thus careful methodological foundations are essential to its success. Over the past few years, some significant methodological advances have occurred in the field, paving the way for more careful work in the future. These include new techniques for the formal analysis of learning algorithms, new experimental approaches to studying learning, successful applications to real-world domains, and the development of integrated cognitive architectures. These changes bode well for this emerging subfield of artificial intelligence. Each of them is discussed briefly in this section.

Theoretical Analyses of Learning Algorithms

Although formal studies of inductive inference have a long history in computer science (Angluin and Smith, 1983), only recently have theorists started to address issues of concern to researchers who actually build machine-learning systems. Initial ideas focused on the notion of convergence: would the learning method eventually construct the exact knowledge desired (Gold, 1967)? This approach constituted an important first step, but it did not afford much insight into realistic learning problems.

A major breakthrough came when researchers turned their focus to the question of evaluating the quality of inductively learned knowledge. The notion of probably approximately correct (PAC) learning forwarded the idea that learned knowledge should usually be relatively accurate when applied in novel situations. Coupling this idea with computational feasibility yields a definition for problems that are polynomially learnable: (1) it must not require too many instances to learn; (2) there must exist an efficient learning method that can produce PAC knowledge; and (3) it must be possible to efficiently determine whether knowledge is consistent with any given instance (Blumer and co-workers, 1987).

When studying the learnability of a particular problem, one common research tactic is to show that one of the three criteria cannot be met. However, because each of these three criteria depends on how acquired knowledge is

represented, results of negative learnability can be brittle. For instance, learning K of N functions using a linear threshold unit with only zero or one weights is not polynomially learnable, but the same class of functions are learnable if the linear threshold unit can use integer weights (Haussler, 1990).

The initial theoretical frameworks focused on learning logical, feature-based concepts in a supervised setting (Kearns and co-workers, 1987), but researchers have since extended the basic framework to other paradigms, including structural concepts (Haussler, 1987), decision lists (Rivest, 1987), conceptual clustering (Pitt and Reinke, 1988), and connectionist networks (Valiant, 1988). They have also addressed learning in the presence of noise (Angluin and Laird, 1988), and they have moved beyond inductive methods to deal with explanation-based methods (Natarajan and Tadepalli, 1988). Even many non-theorists follow this work closely, and many theorists actively read the empirical literature in search of challenging problems.

Experimental Studies of Learning Algorithms

Despite progress on the theoretical front, many learning algorithms remain too complex for formal analysis, and recent progress has also been made in the experimental study of learning methods (Kibler and Langley, 1988). Much of the experimental work has focused on inductive methods (Fahlman, 1988; Fisher, 1987; Quinlan, 1986b; Schlimmer, 1987), but there are also a growing number of experimental studies of explanation-based techniques (Minton and co-workers, 1989; Shavlik, 1990).

One important insight is that performance is the natural dependent measure for such empirical studies, because one can define learning as improvement in performance. There are many measures of performance, including classification accuracy, quality of solution paths, and even CPU time. Different measures are appropriate for different domains and different learning methods, as they may have different goals. However, having at least some measure of performance is essential to evaluating a learning system's behavior. In some cases, intuitively plausible learning methods actually lead to worse performance (Minton, 1985). In other cases, one can use performance measures to determine which components significantly aid the learning process (Schlimmer, 1987).

Researchers have also started to carefully examine the aspects of domains that affect learning behavior. Some experimental studies have focused on naturalistic data in order to show real-world relevance, but others have constructed synthetic domains to allow control of domain characteristics. Two obvious features include complexity of the knowledge to be learned and amount of noise in the data, but others certainly exist. The important point is that many researchers now realize that, in order to make progress, the field requires some explicit methods for evaluating alternative methods and for identifying the conditions under which they work well. Theoretical analyses provide one route to such understanding, but systematic experimentation is another important path.

Common Testbeds and Applications

Early research in machine learning focused on idealized, handcrafted examples, and researchers often tested their systems on only a handful of cases. This has changed drastically in recent years, and papers in the literature now commonly report results on realistic learning tasks that involve many test cases. Moreover, researchers typically report results on a number of different data sets, to show the robustness and generality of their algorithms. The average number of test domains should increase as the standards of the field become higher.

Another encouraging sign is that researchers are starting to test their algorithms on the same task domains, allowing comparisons to be made. This trend has been aided by the collection and distribution of standard data sets. For instance, data on soybean diseases (Michalski and Chilausky, 1980), thyroid diseases (Quinlan, 1987), edibility of mushrooms (Schlimmer, 1987), and Congressional voting records (Fisher, 1987) have been widely distributed and used in testing a number of learning algorithms. Major repositories have emerged, with researchers collecting, documenting, and distributing benchmark data sets. Many of these deal with classification and diagnosis, but standard problem-solving and reasoning tasks are also beginning to emerge.

Despite these encouraging developments, most of these real-world domains remain relatively simple and straightforward. Hopefully, future application efforts will tackle more difficult testbeds that provide greater challenges for machine-learning methods. It is also hoped that the trend will expand to include the documentation and distribution of published algorithms, so that researchers can employ each others' software. This has started to occur within some machine-learning paradigms, but more remains to be done.

Integrated Cognitive Architectures

Another methodological advance relates to the development of integrated architectures for cognition. Early AI researchers commonly implemented a separate system for each new task they encountered. As the field gained experience, high-level languages (eg, production systems) were developed and used to implement new systems, with considerable savings in time and effort. However, these languages incorporated only minor theoretical commitments about the nature of intelligent behavior, and thus provided few constraints on the resulting AI systems. For instance, few formalisms included any automated learning mechanisms.

This trend has changed in recent years, with many researchers now turning to integrated architectures that make strong assumptions about the control structures needed to support intelligence. SOAR (qv) is a prime example of this approach (Laird and co-workers, 1986), and the classifier systems of the genetic algorithm community constitute another instance. Most work in this growing movement includes some automated learning mechanism as an integral part of the architecture, and generality is a central concern, with researchers testing their frameworks on a variety of domains. For instance, the PROD-

IGY (Minton and co-workers, 1989) and THEO (Mitchell and co-workers, in press) architectures incorporate explanation-based methods into their problem-solving engines, whereas ICARUS (Langley and co-workers, in press) relies on case-based concept formation as its main learning mechanism, and DYNA (Sutton, 1990) uses connectionist learning methods. Such integrated frameworks will be necessary if we ever hope to construct intelligent artifacts that can interact with the physical world. We predict that learning will occupy a central role in successful cognitive architectures.

SUMMARY

Over the last decade, the theoretical and methodological advances described in the previous sections have transported machine learning from the sidelines of AI into one of its central foci. Along with this shift has come increased contact with other subcommunities, and as methods for machine learning become more robust, they are gaining increased attention from researchers concerned with planning, diagnosis, natural language, and other problem-oriented areas of artificial intelligence. In turn, these domains provide significant real-world challenges for scientists who have traditionally been concerned with abstract issues in machine learning.

Without doubt, the growing concern with applications will reveal limitations of the existing paradigms and suggest novel directions for automating the acquisition of knowledge. Thus, researchers will be forced to devise new representations, search frameworks, and control schemes to support the learning process. The resulting approaches may initially be domain specific, inefficient, and inelegant, but they will respond to issues that have been previously ignored. Such learning methods may not fit nicely into the organization we have presented, but that is often the nature of scientific progress.

At the same time, others will continue to pursue basic research on learning mechanisms, driven by recognized open issues such as those listed above. These scientists will explore variations and hybrids of existing methods, propose frameworks that unify apparently different techniques, and carry out experimental and theoretical studies to identify the behavior of alternative methods under varying conditions. They should also begin to relate experimental results to those predicted by theory, revising the theory when necessary. Finally, they will attempt to identify new dimensions and new themes that have emerged from the applied work, idealizing them in ways that lets them be studied in the same manner as existing paradigms.

Taken together, basic and applied research in this area should continue to improve the range and capabilities of learning algorithms and to increase our understanding of mechanisms for improving performance with experience. These advances in turn will have far-ranging implications for the rest of artificial intelligence, letting the field move beyond static systems to ones that change their behavior over time as they acquire and refine knowledge.

BIBLIOGRAPHY

- D. H. Ackley, G. E. Hinton, and T. J. Sejnowski, "A Learning Algorithm for Boltzmann Machines," *Cog. Sci.* **9**, 147-169 (1987).
- D. W. Aha and D. Kibler, "Noise-Tolerant Instance-Based Learning Algorithms," *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, Detroit, Morgan-Kaufmann, San Mateo, Calif., 1990, 794-799.
- D. W. Aha, D. Kibler, and M. K. Albert, "Instance-Based Learning Algorithms," *Machine Learning* (in press).
- D. Angluin and P. Laird, "Learning from Noisy Examples," *Machine Learning*, **2**, 343-370 (1988).
- D. Angluin and C. Smith, "Inductive Inference: Theory and Methods," *Comput. Surv.* **15**, 237-269 (1983).
- K. D. Ashley and E. L. Rissland, "Waiting on Weighting: A Symbolic Least Commitment Approach," *Proceedings of the Seventh National Conference on Artificial Intelligence*, St. Paul, Minn., AAAI Press, Menlo Park, Calif., 1988, 239-244.
- É. R. Bareiss, B. W. Porter, and C. C. Wier, "PROTOS: An Exemplar-Based Learning Apprentice," *Proceedings of the Fourth International Workshop on Machine Learning*, Irvine, Calif., Morgan-Kaufmann, San Mateo, Calif., 1987, 12-23.
- R. Barletta and W. Mark, "Explanation-Based Indexing of Cases," *Proceedings of the Seventh National Conference on Artificial Intelligence*, St. Paul, Minn., AAAI Press, Menlo Park, Calif., 1988, 541-546.
- A. G. Barto, R. S. Sutton, and C. W. Anderson, "Neuronlike Elements That Can Solve Difficult Learning Control Problems," *IEEE Trans. System Man Cybernetics*, **13**, 835-846 (1983).
- R. K. Belew and S. Forrest, "Learning and Programming in Classifier Systems," *Machine Learning* **3**, 193-223 (1988).
- A. Blumer, A. Ehrenfeucht, D. Haussler, and M. K. Warmuth, "Occam's Razor," *Info. Proc. Lett.* **24**, 377-380 (1987).
- L. B. Booker, "Classifier Systems That Learn Internal World Models," *Machine Learning* **3**, 161-192 (1988).
- G. Bradshaw, "Learning About Speech Sounds: The NEXUS project," *Proceedings of the Fourth International Workshop on Machine Learning*, Irvine, Calif. Morgan-Kaufmann, 1987, 1-11.
- S. Bradtke and W. G. Lehnert, "Some Experiments with Case-based Search," *Proceedings of the Seventh National Conference on Artificial Intelligence*, St. Paul, Minn., AAAI Press, Menlo Park, Calif. 1988, 133-138.
- M. S. Braverman and R. Wilensky, "Toward a Unification of Case-Based Reasoning and Explanation-Based Learning," *Proceedings of the 1990 AAAI Spring Symposium on Case-Based Reasoning*, Stanford, Calif., 1990.
- L. Brieman, J. H. Friedman, R. A. Olshen, and C. J. Stone, *Classification and Regression Trees*, Wadsworth, Belmont, Calif., 1984.
- J. G. Carbonell, "Derivational Analogy: A Theory of Reconstructive Problem Solving and Expertise Acquisition" in R. S. Michalski, J. G. Carbonell, and T. M. Mitchell, eds., *Machine Learning: An Artificial Intelligence Approach*, vol. 2, Morgan-Kaufmann, San Mateo, Calif., 1986.
- J. G. Carbonell and Y. Gil, "Learning by Experimentation," *Proceedings of the Fourth International Workshop on Machine Learning*, Irvine, Calif., Morgan-Kaufmann, San Mateo, Calif., 1987, 256-266.
- P. Cheeseman, J. Kelly, M. Self, J. Stutz, W. Taylor, and D. Freeman, "AUTOCLASS: A Bayesian Classification System," *Proceedings of the Fifth International Conference on Machine Learning*, Ann Arbor, Mich., Morgan-Kaufmann, San Mateo, Calif., 1988, 54-64.
- P. Clark and T. Niblett, "The CN2 Induction Algorithm," *Machine Learning* **3**, 261-284 (1989).
- T. R. Cover and P. E. Hart, "Nearest Neighbor Pattern Classification," *IEEE Trans. Inf. Theory* **13**, 21-27 (1967).
- R. Davis and D. B. Lenat, *Knowledge-Based Systems in Artificial Intelligence*, McGraw-Hill, New York, 1982.
- G. DeJong and R. J. Mooney, "Explanation-based Learning: An Alternative View," *Machine Learning* **1**, 145-176 (1986)
- T. G. Dietterich, "Learning at the Knowledge Level," *Machine Learning* **1**, 287-316 (1987).
- T. G. Dietterich, H. Hild, and G. Bakiri, "A Comparative Study of ID3 and Backpropagation for English Text-to-Speech Mapping," *Proceedings of the Seventh International Conference on Machine Learning*, Austin, Tex., Morgan-Kaufmann, San Mateo, Calif., 1990, 24-31.
- G. Drastal, G. Czako, and S. Raatz, "Induction in an Abstraction Space: A Form of Constructive Induction," *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, Detroit, Mich., Morgan-Kaufmann, San Mateo, Calif., 1989, 708-712.
- R. Elio and L. Watanabe, "An Incremental Deductive Strategy for Controlling Constructive Induction in Learning from Examples," *Machine Learning* (in press).
- T. Ellman "Approximate Theory Formation: An Explanation-Based Approach," *Proceedings of the Seventh National Conference on Artificial Intelligence*, St. Paul, Minn., AAAI Press, Menlo Park, Calif., 1988, 570-574.
- J. L. Elman, "Finding Structure in Time," *Cog. Sci.* **14**, 179-211 (1990).
- O. Etzioni, "Why PRODIGY/EBL Works," *Proceedings of the Eighth National Conference on Artificial Intelligence*, Boston, AAAI Press, Menlo Park, Calif., 1990, 916-922.
- S. E. Fahlman, "Faster-Learning Variations on Back-Propagation: An Empirical Study," *Proceedings of the 1988 Connectionist Models Summer School*, Pittsburgh, Morgan-Kaufmann, San Mateo, Calif., 1988, 38-51.
- S. E. Fahlman and C. Lebiere, *The Cascade-Correlation Learning Architecture*. Technical Report CMU-CS-90-100, Carnegie Mellon University, School of Computer Science, Pittsburgh, 1990.
- B. Falkenhainer, K. D. Forbus, and D. Gentner, "The Structure-Mapping Engine: Algorithm and Examples," *Artif. Intell.* **41**, 1-63 (1989).
- E. A. Feigenbaum, "The Simulation of Verbal Learning Behavior," in E. A. Feigenbaum and J. Feldman, eds., *Computers and Thought*, McGraw-Hill, New York, 1963.
- R. E. Fikes, P. E. Hart, and N. J. Nilsson, "Learning and Executing Generalized Robot Plans," *Artif. Intell.* **3**, 251-288 (1972).
- D. H. Fisher, "Knowledge Acquisition via Incremental Conceptual Clustering," *Machine Learning* **2**, 139-172 (1987).
- D. H. Fisher, "Noise-Tolerant Conceptual Clustering," *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, Detroit, Morgan-Kaufmann, 1989, 825-830.
- D. H. Fisher and P. Langley, "Approaches to Conceptual Clustering," *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, Los Angeles, Morgan-Kaufmann, San Mateo, Calif., 1985, 691-697).
- D. Fisher and K. B. McKusick, "An Empirical Comparison of ID3 and Back-Propagation," *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, Detroit, Morgan-Kaufmann, San Mateo, Calif., 1989, 788-793.

- D. Geiger, A. Paz, and J. Pearl, "Learning Causal Trees from Dependency Information," *Proceedings of the Eighth National Conference on Artificial Intelligence*, Boston, AAAI Press, Menlo Park, Calif., 1990, 770-776.
- E. Gold, "Language Identification in the Limit," *Inf. Control* 16, 447-474 (1967)
- D. E. Goldberg, "Genetic Algorithms and Rule Learning in Dynamic System Control," *Proceedings of the First International Conference on Genetic Algorithms and Their Applications*, Pittsburgh, Lawrence Erlbaum, Hillsdale, N.J., 1985, 8-15.
- D. E. Goldberg, *Genetic Algorithms in Search, Optimization, and Machine Learning*, Addison-Wesley, Reading, Mass., 1990.
- D. F. Gordon and J. J. Grefenstette, "Explanations of Empirically Derived Reactive Plans," *Proceedings of the Seventh International Conference on Machine Learning*, Austin, Tex., Morgan-Kaufmann, San Mateo, Calif., 1990, 198-203.
- J. J. Grefenstette, C. L. Ramsey, and A. C. Schultz, "Learning Sequential Decision Rules Using Simulation Models and Competition," *Machine Learning* 5, 355-381 (1990).
- K. Hammond, "Learning to Anticipate and Avoid Planning Problems Through the Explanation of Failures," *Proceedings of the Fifth National Conference on Artificial Intelligence*, Philadelphia, AAAI Press, Menlo Park, Calif., 1986, 556-560.
- S. E. Hampson and D. J. Volper, "Disjunctive Models of Boolean Category Learning," *Biol. Cybern.* 56, 121-137 (1987).
- D. Haussler, "Bias, Version Spaces, and Valiant's Learning Framework," *Proceedings of the Fourth International Workshop on Machine Learning*, Irvine, Calif., Morgan-Kaufmann, San Mateo, Calif., 1987, 324-336.
- D. Haussler, "Probably Approximately Correct Learning," *Proceedings of the Eighth National Conference on Artificial Intelligence*, Boston, AAAI Press, Menlo Park, Calif., 1990, 1101-1108.
- G. E. Hinton, "Learning Distributed Representations of Concepts," *Proceedings of the Eighth Annual Conference of the Cognitive Science Society*, Amherst, Mass., Lawrence Erlbaum, Hillsdale, N.J., 1986, 1-12.
- G. E. Hinton, "Connectionist Learning Procedures," *Artif. Intel.* 40, 185-234 (1989).
- H. Hirsh, "Combining Empirical and Analytical Learning with Version Spaces," *Proceedings of the Sixth International Workshop on Machine Learning*, Ithaca, N.Y., Morgan-Kaufmann, Publishers, San Mateo, Calif., 1989.
- J. H. Holland, "Properties of the Bucket Brigade Algorithm," *Proceedings of the First International Conference on Genetic Algorithms and Their Applications*, Pittsburgh, Lawrence Erlbaum, Hillsdale, N.J., 1985, 1-7.
- L. Hunter, "Planning to Learn," *Proceedings of the Twelfth Annual Conference of the Cognitive Science Society* Cambridge, Mass., Lawrence Erlbaum, Hillsdale, N.J., 1990, 261-268.
- G. A. Iba, "A Heuristic Approach to the Discovery of Macro-Operators," *Machine Learning* 3, 285-317 (1989).
- R. Jones, *A Model of Retrieval in Problem Solving*, Technical Report 89-27, Ph.D. dissertation, University of California, Irvine, Department of Information and Computer Science, 1989.
- M. Kearns, M. Li, L. Pitt, and L. G. Valiant, "Recent Results on Boolean Concept Learning," *Proceedings of the Fourth International Workshop on Machine Learning*, Irvine, Calif., Morgan-Kaufmann, San Mateo, Calif., 1987, 337-352.
- R. M. Keller, "Concept Learning in Context," *Proceedings of the Fourth International Workshop on Machine Learning*, Irvine, Calif., Morgan-Kaufmann, San Mateo, Calif., 1987, 91-102.
- D. Kibler and P. Langley, "Machine Learning as an Experimental Science," *Proceedings of the Third European Working Session on Learning*, Pittman, London, 1988.
- K. Knight, *A Gentle Introduction to Subsymbolic Computation: Connectionist for the A.I. Researcher*, Technical Report CMU-CS-89-150, Carnegie Mellon University, School of Computer Science, Pittsburgh, 1989.
- C. Knoblock, "Learning Abstraction Hierarchies for Problem Solving," *Proceedings of the Eighth National Conference on Artificial Intelligence*, Boston, AAAI Press, Menlo Park, Calif., 1990, 923-928.
- T. Kohonen, E. Oja, and P. Lehtiö, "Storage and Processing of Information in Distributed Associative Memory Systems," in G. E. Hinton and J. A. Anderson, eds., *Parallel Models of Associative Memory*, Lawrence Erlbaum, Hillsdale, N.J., 1981.
- J. L. Kolodner, "Maintaining Organization in a Dynamic Long-term Memory," *Cog. Sci.* 7, 243-280 (1983).
- J. L. Kolodner, "Selecting the Best Case for a Case-Based Reasoner," *Proceedings of the Eleventh Annual Conference of the Cognitive Science Society*, Ann Arbor, Mich., Lawrence Erlbaum, Hillsdale, N.J., 1989, 155-162.
- I. Kononenko, I. Bratko, and E. Roskar, *Experiments in Automatic Learning of Medical Diagnostic Rules*, Josef Stefan Institute, Ljubljana, Yugoslavia, 1984.
- P. Koton, "Integrating Case-based and Causal Reasoning," *Proceedings of the Tenth Annual Conference of the Cognitive Science Society*, Montreal, Quebec, Canada, Lawrence Erlbaum, Hillsdale, N.J., 1988a, 167-173.
- P. Koton, "Reasoning About Evidence in Causal Explanation," *Proceedings of the Seventh National Conference on Artificial Intelligence*, St. Paul, Minn., AAAI Press, Menlo Park, Calif., 1988b, 256-261.
- J. R. Koza, "Hierarchical Genetic Algorithms Operating on Populations of Computer Programs," *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, Detroit, Morgan-Kaufmann, San Mateo, Calif., 1989, 768-774.
- J. E. Laird, "Recovery from Incorrect Knowledge in SOAR," *Proceedings of the Seventh National Conference on Artificial Intelligence*, St. Paul, Minn., AAAI Press, Menlo Park, Calif., 1988, 618-623.
- J. E. Laird, P. S. Rosenbloom, and A. Newell, "Chunking in SOAR: The Anatomy of a General Learning Mechanism," *Machine Learning* 1, 11-46 (1986).
- P. Langley, "Learning to Search: From Weak Methods to Domain-Specific Heuristics," *Cog. Sci.* 9, 217-260 (1985).
- P. Langley, "Unifying Themes in Empirical and Explanation-based Learning," *Proceedings of the Sixth International Workshop on Machine Learning*, Ithaca, N.Y., Morgan-Kaufmann, San Mateo, Calif., 1989.
- P. Langley, J. H. Gennari, and W. Iba, "Hill-Climbing Theories of Learning," *Proceedings of the Fourth International Workshop on Machine Learning*, Irvine, Calif., Morgan-Kaufmann, San Mateo, Calif., 1987, 312-323.
- P. Langley, K. Thompson, W. Iba, J. H. Gennari, and J. A. Allen, "An Integrated Cognitive Architecture for Autonomous Agents," in W. Van De Velde, ed., *Representation and Learning in Autonomous Agents*, North Holland, Amsterdam (in press).
- M. Lebowitz, "Experiments with Incremental Concept Formation: UNIMEM," *Machine Learning* 2, 103-138 (1987).
- S. Marcus, ed., *Automating Knowledge Acquisition for Expert Systems*, Kluwer, Boston, 1988.
- S. Markovitch and P. D. Scott, "Utilization Filtering: A Method

- for Reducing the Inherent Harmfulness of Deductively Learned Knowledge," *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, Detroit, Morgan-Kaufmann, San Mateo, Calif., 1989, 738-743.
- C. J. Matheus and L. A. Rendell, "Constructive Induction on Decision Trees," *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, Detroit, Morgan-Kaufmann, San Mateo, Calif., 1989, 645-650.
- R. S. Michalski, "A Theory and Methodology of Learning from Examples," in R. S. Michalski, J. G. Carbonell, and T. M. Mitchell, eds., *Machine Learning: An Artificial Intelligence Approach*, vol. 1, Morgan-Kaufmann, San Mateo, Calif., 1983.
- R. S. Michalski, "How to Learn Imprecise Concepts: A Method for Employing a Two-Tiered Knowledge Representation in Learning," *Proceedings of the Fourth International Workshop on Machine Learning*, Irvine, Calif., Morgan-Kaufmann, San Mateo, Calif., 1987, 50-58.
- R. S. Michalski and R. L. Chilausky, "Learning by Being Told and Learning from Examples: An Experimental Comparison of Two Methods of Knowledge Acquisition in the Context of Developing an Expert System for Soybean Disease Diagnosis," *Int. J. Policy Anal. Infor. Syst.* 4 (1980).
- R. S. Michalski and R. Stepp, "Learning from Observation: Conceptual Clustering," in R. S. Michalski, J. G. Carbonell, and T. M. Mitchell, eds., *Machine Learning: An Artificial Intelligence Approach*, vol. 1, Morgan-Kaufmann, San Mateo, Calif., 1983.
- M. Minsky and S. Papert, *Perceptrons: An Introduction to Computational Geometry*, MIT Press, Cambridge, Mass., 1969.
- S. N. Minton, "Selectively Generalizing Plans for Problem Solving," *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, Morgan-Kaufmann, San Mateo, Calif., 1985, 596-599.
- S. N. Minton, "Quantitative Results Concerning the Utility of Explanation-Based Learning," *Artif. Intell.* 42, 363-391 (1990).
- S. N. Minton, J. G. Carbonell, C. A. Knoblock, D. R. Kuokka, O. Etzioni, and Y. Gil, "Explanation-Based Learning: A Problem Solving Perspective," *Artif. Intell.* 40, 63-118 (1989).
- T. M. Mitchell, "Version Spaces: A Candidate Elimination Approach to Rule Learning," *Proceedings of the Fifth International Joint Conference on Artificial Intelligence*, Cambridge, Mass., Morgan-Kaufmann, San Mateo, Calif., 1977, 305-310.
- T. M. Mitchell, J. Allen, P. Chalasani, J. Cheng, O. Etzioni, M. Ringuette, and J. C. Schlimmer, "THEO: A Framework for Self-Improving Systems," in K. VanLehn, ed., *Architectures for Intelligence*, Lawrence Erlbaum, Hillsdale, N.J., (in press).
- T. M. Mitchell, R. M. Keller, and S. T. Kedar-Cabelli, "Explanation-Based Generalization: A Unifying View," *Machine Learning* 1, 47-80 (1986).
- T. M. Mitchell, S. Mahadevan, and L. Steinberg, "LEAP: A Learning Apprentice for VLSI Design," *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, Los Angeles, Morgan-Kaufmann, San Mateo, Calif., 1985, 573-580.
- T. M. Mitchell, P. E. Utgoff, and R. B. Banerji, "Learning Problem Solving Heuristics by Experimentation," in R. S. Michalski, J. G. Carbonell, and T. M. Mitchell, eds., *Machine Learning: An Artificial Intelligence Approach*, vol. 1, Morgan-Kaufmann, San Mateo, Calif., 1983.
- R. J. Mooney, J. W. Shavlik, G. G. Towell, and A. Gove, "An Experimental Comparison of Symbolic and Connectionist Learning Algorithms," *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, Detroit, Morgan-Kaufmann, San Mateo, Calif., 1990.
- J. Mostow, "Design by Derivational Analogy: Issues in the Automated Replay of Design Plans," *Artif. Intell.* 40, 119-184 (1989).
- M. Mozer and J. Bachrach, "SLUG: A Connectionist Architecture for Inferring the Structure of Finite-State Environments," *Machine Learning* (in press).
- S. Muggleton, and W. Buntine, "Machine Invention of First-Order Predictates by Inverting Resolution," *Proceedings of the Fifth International Conference on Machine Learning*, Ann Arbor, Mich., Morgan-Kaufmann, San Mateo, Calif., 1988, 339-352.
- B. K. Natarajan and P. Tadepalli, "Two New Frameworks for Learning," *Proceedings of the Fifth International Conference on Machine Learning*, Ann Arbor, Mich., Morgan-Kaufmann, San Mateo, Calif., 1988, 402-415.
- H. T. Ng and R. J. Mooney, "On the Role of Coherence in Abductive Explanation," *Proceedings of the Eighth National Conference on Artificial Intelligence*, Boston, 1990, 337-342.
- N. Nilsson, *Learning Machines*. McGraw-Hill, New York, 1965.
- D. Ourston and R. J. Mooney, "Changing the Rules: A Comprehensive Approach to Theory Refinement," *Proceedings of the Eighth National Conference on Artificial Intelligence*, Boston, AAAI Press, Menlo Park, Calif., 1990, 815-820.
- C. Owens, "Integrating Feature Extraction and Memory Search," *Proceedings of the Eleventh Annual Conference of the Cognitive Science Society*, Ann Arbor, Mich., Lawrence Erlbaum, Hillsdale, N.J., 1989, 163-170.
- G. Pagallo, "Learning DNF by Decision Trees," *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, Detroit, Morgan Kaufmann, San Mateo, Calif., 1989, 639-644.
- L. Pitt and R. E. Reinke, "Criteria for Polynomial-Time (Conceptual) Clustering," *Machine Learning* 2, 371-396 (1988).
- R. Prager, T. D. Harrison, and F. Fallside, "Boltzmann Machines for Speech Recognition," *Comput. Speech Lang.* 1, 1-20 (1986).
- J. R. Quinlan, "Learning Efficient Classification Procedures and Their Application to Chess Endgames," in R. S. Michalski, J. G. Carbonell, and T. M. Mitchell, eds., *Machine Learning: An Artificial Intelligence Approach*, vol. 1, Morgan-Kaufmann, San Mateo, Calif., 1983.
- J. R. Quinlan, "Induction of Decision Trees," *Machine Learning* 1, 81-106 (1986a).
- J. R. Quinlan, "The Effect of Noise on Concept Learning," in R. S. Michalski, J. G. Carbonell, and T. M. Mitchell, eds., *Machine Learning: An Artificial Intelligence Approach*, vol. 2, Morgan-Kaufmann, San Mateo, Calif., 1986b.
- J. R. Quinlan, "Generating Production Rules from Decision Trees," *Proceedings of the Tenth International Joint Conference on Artificial Intelligence*, Milan, Italy, Morgan-Kaufmann, 1987, 304-307.
- J. R. Quinlan, P. J. Compton, K. A. Horn, and L. Lazarus, "Inductive Knowledge Acquisition: A Case study," *Proceedings of the Second Australian Conference on Applications of Expert Systems*, Sydney, Australia, 1986.
- M. Redmond, "Combining Case-based Reasoning, Explanation-based Learning, and Learning from Instruction," *Proceedings of the Sixth International Workshop on Machine Learning*, Ithaca, N.Y., Morgan-Kaufmann, San Mateo, Calif., 1989, 20-22.
- M. Redmond, "Distributed Cases for Case-based Reasoning: Facilitating Use of Multiple Cases," *Proceedings of the Eighth National Conference on Artificial Intelligence*, Boston, AAAI Press, Menlo Park, Calif., 1990, 304-309.

- R. L. Rivest, "Learning Decision Lists," *Machine Learning* 2, 229-246 (1987).
- G. G. Robertson, "Population Size in Classifier Systems," *Proceedings of the Fifth International Conference on Machine Learning*, Ann Arbor, Mich., Morgan-Kaufmann, San Mateo, Calif., 1988, 142-152.
- F. Rosenblatt, *Principles of Neurodynamics*, Spartan Books, New York, 1962.
- P. S. Rosenbloom, J. E. Laird, J. McDermott, A. Newell, and E. Orciuch, "R1-SOAR: An Experiment in Knowledge-Intensive Programming in a Problem-Solving Architecture," *IEEE Trans. Pattern Anal. Mach. Intell.* 7, 561-569 (1985).
- D. E. Rumelhart and J. L. McClelland, "On Learning the Past Tenses of English Verbs," in J. L. McClelland and D. E. Rumelhart, eds., *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, vol. 2, MIT Press, Cambridge, Mass., 1986.
- S. Salzberg, *Learning with Nested Generalized Exemplars*, Kluwer, Boston, 1990.
- A. L. Samuel, "Some Studies in Machine Learning Using the Game of Checkers," *IBM J. Res. Dev.* 3, 210-229 (1959).
- J. D. Schaffer and J. J. Grefenstette, "Multi-Objective Learning via Genetic Algorithms," *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, Los Angeles, Morgan-Kaufmann, San Mateo, Calif., 1985, 593-595.
- J. C. Schlimmer, "Incremental Adjustment of Representations for Learning," *Proceedings of the Fourth International Workshop on Machine Learning*, Irvine, Calif., Morgan-Kaufmann, San Mateo, Calif., 1987, 79-90.
- J. C. Schlimmer and D. H. Fisher, "A Case Study of Incremental Concept Induction," *Proceedings of the Fifth National Conference on Artificial Intelligence*, Philadelphia, 1986, 496-501.
- J. C. Schlimmer and R. H. Granger, Jr., "Beyond Incremental Processing: Tracking Concept Drift," *Proceedings of the Fifth National Conference on Artificial Intelligence*, Philadelphia, 1986, 502-507.
- T. J. Sejnowski and C. R. Rosenberg, "Parallel Networks that Learn to Pronounce English Text," *Complex Systems* 1, 145-168 (1987).
- C. G. Shaefer, "The ARGOT Strategy: Adaptive Representation Genetic Optimizer Technique," *Proceedings of the Second International Conference on Genetic Algorithms*, 1987.
- J. W. Shavlik, (1990). "Acquiring Recursive and Iterative Concepts with Explanation-based Learning," *Machine Learning*, 5, 39-70 (1990).
- D. Sleeman, P. Langley, and T. Mitchell, "Learning from Solution Paths: An Approach to the Credit Assignment Problem," *AI Mag.* 3, 48-52 (1982).
- R. G. Smith, H. A. Winston, T. M. Mitchell, and B. G. Buchanan, "Representation and Use of Explicit Justifications for Knowledge Base Refinement," *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, Los Angeles, Morgan-Kaufmann, Publishers, San Mateo, Calif., 1985, 673-680.
- S. F. Smith, "Flexible Learning of Problem Solving Heuristics Through Adaptive Search," *Proceedings of the Eighth International Joint Conference on Artificial Intelligence*, Karlsruhe, Germany, Morgan-Kaufmann, Publishers, San Mateo, Calif., 1983, 422-425.
- C. W. Stanfill, "Memory-Based Reasoning Applied to English Pronunciation," *Proceedings of the Sixth National Conference on Artificial Intelligence*, Seattle, AAAI Press, Menlo Park, Calif., 1987, 577-581.
- R. E. Stepp and R. S. Michalski, "Conceptual Clustering: Inventing Goal-Oriented Classifications of Structured Objects," in R. S. Michalski, J. G. Carbonell, and T. M. Mitchell, eds., *Machine Learning: An Artificial Intelligence Approach*, vol. 2, Morgan-Kaufmann, San Mateo, Calif., 1986.
- R. S. Sutton, "Integrated Architectures for Learning, Planning, and Reacting based on Approximating Dynamic Programming," *Proceedings of the Seventh International Conference on Machine Learning*, Austin, Tex., Morgan-Kaufmann, San Mateo, Calif., 1990, 216-224.
- K. Sycara, "Patching Up Old Plans," *Proceedings of the Tenth Annual Conference of the Cognitive Science Society*, Montreal, Quebec, Canada, Lawrence Erlbaum, Hillsdale, N.J., 1988, 405-411.
- P. Tadepalli, "Planning in Games Using Approximately Learned Macros," *Proceedings of the Sixth International Workshop on Machine Learning*, Ithaca, N.Y., Morgan-Kaufmann, San Mateo, Calif., 1989, 221-223.
- M. Tambe and P. Rosenbloom, "Eliminating Expensive Chunks by Restricting Expressiveness," *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, Detroit, Morgan-Kaufmann, San Mateo, Calif., 1989 731-737.
- G. G. Towell, J. W. Shavlik, and M. O. Noordewier, "Refinement of Approximate Domain Theories by Knowledge-based Neural Networks," *Proceedings of the Eighth National Conference on Artificial Intelligence*, Boston, AAAI Press, Menlo Park, Calif., 1990, 861-866.
- P. E. Utgoff, "Perceptron Trees: A Case Study in Hybrid Concept Representations," *Proceedings of the Seventh National Conference on Artificial Intelligence*, St. Paul, Minn., AAAI Press, Menlo Park, Calif., 1988, 601-606.
- P. E. Utgoff, "Incremental Induction of Decision Trees," *Machine Learning* 4, 161-186 (1989).
- L. G. Valiant, "Functionality in Neural Networks," *Proceedings of the Seventh National Conference on Artificial Intelligence*, St. Paul, Minn., AAAI Press, Menlo Park, Calif., 1988, 629-640.
- L. Watanabe and L. Rendell, "Effective Generalization of Relational Description," *Proceedings of the Eighth National Conference on Artificial Intelligence*, Boston, AAAI Press, Menlo Park, Calif., 1990, 875-881.
- S. M. Weiss and I. Kapouleas, "An Empirical Comparison of Pattern Recognition, Neural Nets, and Machine Learning Classification Methods," *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, Detroit, Morgan-Kaufmann, 1989, 781-787.
- B. Widrow and M. E. Hoff, "Adaptive Switching Circuits. IRE WESCON," *Convention Record* 4 (1960).
- S. W. Wilson, "Classifier Systems and the Animat Problem," *Machine Learning* 2, 199-228 (1987).
- P. H. Winston, "Learning New Principles from Precedents and Exercises," *Artif. Intell.* 19, 321-350 (1984).

JEFFREY C. SCHLIMMER
Carnegie Mellon University
PAT LANGLEY
NASA Ames Research Center

The authors appreciate the thoughtful comments and suggestions of Andrew Barto, Gerald DeJong, John Grefenstette, Janet Kolodner, Amy Lansky, and Edwina Rissland. Any errors or ambiguities are the sole responsibility of the authors.

ENCYCLOPEDIA OF ARTIFICIAL INTELLIGENCE

This extensively revised and expanded *Second Edition* of the *Encyclopedia of Artificial Intelligence* defines the discipline by bringing together the core of knowledge from all fields encompassed by AI. It covers the latest developments in current AI topics such as neural networks, fuzzy logic, machine vision, natural language generation, and many more. Includes:

- Over 450 articles—all entries written expressly for the *Encyclopedia*
- Over 5,000 literature references; 454 illustrations and color photographs
- Over 50% new and revised material
- Exemplary indexing and cross-referencing for easy, complete information access to all topics

Praise for the *First Edition*...

"The *Encyclopedia* is a wonder of clarity and scope: surprisingly easy to read...the clarity is an especially pleasant surprise, considering the articles were all written by AI experts...It's a treasure house of easily accessible knowledge."

—*Language Technology*

"Excellent bibliographies are attached to most of the articles, and diagrams and sketches are clear and helpful. The indexing and cross-indexing are exemplary. As the editor points out, the reader will be led by the extensive cross-references to almost every other article..."

—*Artificial Intelligence Reporter*

"The *Encyclopedia* is a first-class piece of work that will be an indispensable part of any AI library..."

—*Computing Reviews*

"...A tour de force...a truly fantastic encyclopedia which no one in the field of artificial intelligence can afford to be without."

—*Systems Research & Information*

WILEY-INTERSCIENCE

John Wiley & Sons, Inc.
Professional, Reference and Trade Group
605 Third Avenue, New York, NY 10158-0012
New York • Chichester • Brisbane • Toronto • Singapore
ISBN 0 471-50307-X (Two-Volume Set)
ISBN 0 471-50305-3 (Vol. 1)
ISBN 0 471-50306-1 (Vol. 2)

ISBN 0-471-50305-3



9 780471 503057